

IDA

INSTITUTE FOR DEFENSE ANALYSES

Quantum Algorithms and Protocols

Steve Huntsman

February 2001

Approved for public release;
distribution unlimited.

IDA Paper P-3540

Log: H 00-001730

20010328 167

This work was conducted under IDA's independent research program, CRP-2048. The publication of this IDA document does not indicate endorsement by the Department of Defense, nor should the contents be construed as reflecting the official position of that Agency.

© 2000, 2001 Institute for Defense Analyses, 1801 N. Beauregard Street, Alexandria, Virginia 22311-1772 • (703) 845-2000.

This material may be reproduced by or for the U.S. Government.

INSTITUTE FOR DEFENSE ANALYSES

IDA Paper P-3540

Quantum Algorithms and Protocols

Steve Huntsman

PREFACE

This work was performed under an IDA Central Research Project entitled
“Quantum Computing.”

ACKNOWLEDGMENTS

The author wishes to thank Jim Heagy and Mel Currie for reviewing the text, offering improvements (in particular with respect to the material on decoherence presented herein) and catching several errors. Bohdan Balko, Bobby Doney, Jon Dowling, Mark Heiligman, Seth Lloyd, and Jeffrey Yezpe also gave helpful insights. Finally, the author wishes to thank Bob Clancy for a fruitful discussion concerning monkeys wearing Polaroid sunglasses (as opposed to hammering at typewriters).

CONTENTS

I. INTRODUCTION: THE BASIS OF QUANTUM COMPUTING	I-1
A. Schrödinger, EPR, and Bell (or, When a Tree Falls in the Forest...)	I-1
B. Prototype Two-State Quantum Systems: Qubits.....	I-1
C. Quantum Parallelism via Entanglement	I-2
D. Decoherence.....	I-3
E. The State of the Art	I-5
II. BASIC THEORETICAL MODELS	II-1
A. Logic Gates	II-1
B. Feynman's Quantum Metaprogram.....	II-2
C. The Quantum Fourier Transform	II-2
D. Quantum Algorithms for Special Oracle Problems.....	II-3
III. FACTORING ON A QUANTUM COMPUTER.....	III-1
A. Shor's Algorithm.....	III-1
B. Factoring as an Instance of the Abelian Stabilizer Problem	III-2
IV. SEARCH ALGORITHMS	IV-1
A. Finding a Needle in a Quantum Haystack	IV-1
B. Generalized Unstructured and Unstructured Parallel Quantum Searching.....	IV-2
C. Structured Quantum Searching	IV-3
V. OTHER SELECTED ALGORITHMS.....	V-1
A. Integration	V-1
B. Simulation of Local Quantum Systems	V-2
C. Quantum Cellular Automata	V-4

VI. QUANTUM INFORMATION THEORY	VI-1
A. Quantum Communication Channels.....	VI-1
B. Quantum Error-Correcting Codes	VI-3
C. Quantum Key Distribution.....	VI-6
D. Exploiting Entanglement: Quantum Teleportation and Communication Complexity	VI-8
VII. CONCLUSION	VII-1
APPENDIX A—Hard Number—Theoretic Problems	A-1
APPENDIX B—Classical Information Theory.....	B-1
APPENDIX C—References.....	C-1

I. INTRODUCTION: THE BASIS OF QUANTUM COMPUTING

A. SCHRÖDINGER, EPR, AND BELL (OR, WHEN A TREE FALLS IN THE FOREST...)

Quantum information is a major initiative in the physical and informational sciences which traces its roots back to the *gedanken* experiments of Schrödinger and Einstein, Podolsky, and Rosen (EPR). EPR, following a “Schrödinger’s cat” line of thinking in an attempt to validate the “no-dice” opposition to quantum theory, pointed out that the linear superposition principle of quantum mechanics implied that so-called “entangled” states allowed by the theory could be created in such a way as to violate supposedly natural criteria (such as locality). In their quest for a deterministic explanation, EPR concluded that quantum mechanics was invalid and that the notional collapse of a superposition of quantum states was illusory. That is, the evolution of quantum states was deterministic but somehow hidden from measurement. Bell [12] sought to address these ideas in formulating the Bell inequalities for hidden variable theories; these inequalities were experimentally testable hypotheses which could conclusively confirm or deny the nondeterminism of quantum mechanics.

The first test of Bell’s inequalities was conducted by Laméhi-Rachti and Mittig [90] in 1976; their results disagreed with hidden-variable interpretations but were inconclusive. Later experiments progressively wore away at hidden-variable theories; in 1996, the creation of a quantum superposition was experimentally verified [106], and the days of the hidden-variable theories were conclusively over.

The resolution of this basically metaphysical issue has significant implications for national security and the physical and informational sciences. Its potentially profound effect on the evolution of these fields (and, in particular, their intersection) is the motivation for this discussion.

B. PROTOTYPE TWO-STATE QUANTUM SYSTEMS: QUBITS

Consider (via identification or analogy with, e.g., the polarization of a photon or the spin of an electron) a quantum-mechanical system whose Hilbert or state space \mathcal{H} is generated by two basis states, denoted $|0\rangle$ and $|1\rangle$. A general state $|a\rangle$ is then a unit norm

linear superposition of the basis states. That is, we have $|a\rangle = \alpha|0\rangle + \beta|1\rangle$, with $|\alpha|^2 + |\beta|^2 = 1$. We refer to such a state as a *qubit*. The probability that measurement of a qubit $|a\rangle$ results in the outcome of state $|0\rangle$ (resp., $|1\rangle$) is $|\alpha|^2$ (resp., $|\beta|^2$). We can use an operator rather than a state vector to describe the system; in this setting we consider the *density matrix* $\rho_a \equiv |a\rangle\langle a|$. Measurements correspond to the projection operators $\pi_0 \equiv |0\rangle\langle 0|$ (resp., $\pi_1 \equiv |1\rangle\langle 1|$), and the associated probabilities can be obtained by noting that $\text{Tr}(|0\rangle\langle 0|\rho_a) = \alpha$ (resp., $\text{Tr}(|1\rangle\langle 1|\rho_a) = \beta$).

Now consider a collection of n such systems: the Hilbert space \mathcal{H}^n of the resulting composite system is the tensor product of the subsystems, which has dimension 2^n . A general state is now a linear superposition of the basis states (which can be expressed as bit vectors or decimal numbers in the canonical computational basis) in \mathcal{H}^n . It is easy to see that there are then necessarily states in \mathcal{H}^n which are not themselves tensor products of qubits; these are referred to as *entangled states*. Using the shorthand $|ab\rangle$ or $|a\rangle|b\rangle$ for $|a\rangle \otimes |b\rangle$ (and ignoring a normalization factor), we note, for example, that $|00\rangle + |11\rangle$ is such a (maximally) entangled state, called an *EPR pair*, of which a measurement in the computational basis can result in only two possible outcomes ($|00\rangle$ or $|11\rangle$)—whereas in general a measurement of a two-qubit system may result in any of four possible outcomes ($|00\rangle$, $|01\rangle$, $|10\rangle$ or $|11\rangle$).

C. QUANTUM PARALLELISM VIA ENTANGLEMENT

The notorious difficulty of the quantum-mechanical N -body problem is a consequence of the fact that linear growth in the number of particles results in exponential growth in the dimension of the Hilbert space—and hence in the cost of simulation. Theoretical work on the thermodynamics of classical computation by Bennett [15] and Fredkin and Toffoli [60] and on the simulation of Turing machines with quantum systems by Benioff [13] led Feynman [57] to argue that this problem could in some sense be its own solution: a quantum mechanical system that was more or less impossible to simulate classically could be effectively simulated by another quantum mechanical system. (In fact, a quantum computer can do the job, as we will see; quantum mechanical simulation would probably be the first real use of the technology if a fully operational quantum computer with over 30 or so qubits was developed [94], [95].)

Introducing the *Walsh-Hadamard operator* (with matrix in the computational basis),

$$W = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix},$$

we see that $W|0\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$. That is, applying the Walsh-Hadamard operator to the “ground” state gives a uniform superposition of the basis states. This operator (geometrically realized as the composition of a rotation and a reflection) is a precursor to more sophisticated unitary operators or *quantum gates*. If we consider the tensor product W_n (acting on \mathcal{H}^n) of n single-qubit Walsh-Hadamard operators, we obtain

$$W_n|i\rangle \equiv \bigotimes_{k=1}^n W_{(k)}|i_k\rangle = \frac{1}{\sqrt{2^n}} \sum_{j=0}^{2^n-1} (-1)^{(i,j)} |j\rangle$$

where $i \equiv \sum_{k=1}^n i_k 2^k$, $|i\rangle \equiv \bigotimes_{k=1}^n |i_k\rangle$; $(i,j) \equiv \sum_{k=1}^n i_k j_k$.

Forming $W_n|0\dots 0\rangle = W|0\rangle \otimes \dots \otimes W|0\rangle$, we obtain a uniform superposition of all the basis states in the total Hilbert space; a measurement results in an outcome of an arbitrary bit string of length n with probability 2^{-n} . Applying a quantum gate to this superposition is equivalent to superposing the states resulting from applying the gate to each (suitably normalized) basis state. This is the prototype of *quantum parallelism*.

Although measurements of such a state give a procedure for generating random numbers, it is far from clear how to generalize it (much less actually physically implement it) in such a way as to actually do anything useful that could not be achieved much more easily with (say) a radioactive decay source. Indeed, the bright light of quantum parallelism casts a dark shadow of quantum measurement and decoherence. Even if we can somehow implement a technique for entangling and manipulating qubits, we are lost without a way to measure the desired basis state with a probability greater than 2^{-n} or if the environment collapses our superpositions. The discovery of a realizable algorithmic technique (the quantum Fourier transform) by Coppersmith [44] for generating constructive interference of desirable states marked a crucial step towards realizing the utility of quantum computation.

D. DECOHERENCE

One of the fundamental tenets of quantum mechanics is that a measurement collapses a quantum superposition into a fixed state. The question of precisely what defines a measurement is subtle, however; indeed, interactions between the external environment and a quantum superposition will generally force nondiagonal elements of

the density matrix to become negligible, and the initially coherent phases of subsystems will decohere.

Efficient algorithms for quantum computation can provide an answer to the measurement problem as applied to a system of qubits isolated from their environment, but in fact a system of qubits will rapidly interact with its environment—no matter how weak the coupling—and the superposition will effectively collapse. The degree to which this process of *decoherence* can be delayed is the x-factor in building real-world quantum computers.

Omnès [109] sketches a mechanism for the decoherence of a single qubit as a consequence of interaction with an (internal) environment of n (externally) noninteracting qubits; we will follow his treatment. Such a system can be described by a Hamiltonian of the form

$$H = H_{\text{int}} = \sigma \sum_{k=1}^n g_k \bigotimes_{j=1}^{k-1} Id_j \otimes \sigma_k \otimes \bigotimes_{j=k+1}^n Id_j ,$$

where $\sigma_{(k)} = -|0\rangle\langle 0|_{(k)} + |1\rangle\langle 1|_{(k)}$, $Id_{(k)}$ denote single-qubit identity operators, and g_k are coupling constants. The state

$$|\Psi(t)\rangle = a|0\rangle \bigotimes_{k=1}^n (\alpha_k e^{ig_k t} |0\rangle_k + \beta_k e^{-ig_k t} |1\rangle_k) + b|1\rangle \bigotimes_{k=1}^n (\alpha_k e^{-ig_k t} |0\rangle_k + \beta_k e^{ig_k t} |1\rangle_k)$$

then satisfies the Schrödinger equation; the reduced density matrix for the qubit (obtained by performing a partial trace over the environmental degrees of freedom) is

$$\rho = \text{Tr}_k |\Psi(t)\rangle\langle\Psi(t)| = |a|^2 |0\rangle\langle 0| + |b|^2 |1\rangle\langle 1| + z(t) a \bar{b} |0\rangle\langle 1| + \bar{z}(t) \bar{a} b |1\rangle\langle 0| ;$$

$$z(t) \equiv \prod_{k=1}^n \left[\cos 2g_k t + i(|\beta_k|^2 - |\alpha_k|^2) \sin 2g_k t \right] .$$

It can be shown that

$$\langle |z(t)|^2 \rangle = 2^{-n} \prod_{k=1}^n \left[1 + (|\beta_k|^2 - |\alpha_k|^2)^2 \right] .$$

If the initial state of the environment (i.e., the distribution of α_k, β_k) is random, then this quantity is exponentially small. This effective diagonalisation of the density matrix is the hallmark of decoherence: the probabilities of quantum superpositions decrease rapidly as a result of interactions.

Avoiding decoherence in experiments (much less actual physical quantum computers or channels) is made especially difficult because of interactions with the

external environment, which is a much harder problem to address than avoiding undesired qubit-qubit interactions. Consider for example the case of a harmonic oscillator weakly coupled to a bath of harmonic oscillators [109]. If the harmonic oscillator is prepared in a superposition of two harmonic oscillator coherent states, the decay of the off-diagonal elements is exponential with a characteristic time given by

$$\tau_0 = \frac{2\hbar\tau}{m\omega(x_1(0) - x_2(0))^2} ,$$

(the expression given in [109] contains errors) where τ is the damping time of the oscillator. For a quartz oscillator with fundamental frequency $f = 2\pi\omega = 50$ MHz, mass $m \approx 10^{-25}$ kg (mass of the SiO_2 molecule), initial coherent state separation $\Delta x = x_1(0) - x_2(0) = 10^{-10}$ m (1 Å), and $Q = \omega\tau \approx 10^3$, we find $\tau_0 \approx 3$ s. If we include the effects of temperature the characteristic time is [137]

$$\tau_T = \frac{\hbar^2\tau}{2mkT(x_1(0) - x_2(0))^2} .$$

The ratio of characteristic times is given by $\tau_T/\tau_0 = \hbar\omega/kT$, that is, by the ratio of excited to thermal energies. So the same quartz oscillator at $T = 300$ K has a characteristic decoherence time of about 0.3 μs , indicating the dramatic and constraining effects of temperature. (DiVincenzo [49] lists decoherence times for other physical systems that have been proposed for quantum computer realizations.). Low temperatures can delay the onset of decoherence (indeed, the first evidence of a macroscopic quantum superposition was recently obtained for a superconducting quantum interference device at a few degrees Kelvin [A1]).

To use a system, it must have some coupling to the external environment, and it is therefore subject to rapid decoherence. This Catch-22 can be circumvented by employing quantum error-correcting codes, whose independent discovery by Shor and Calderbank [34] and Steane [127] made quantum information technology a realistic goal. But whether that goal will ever be fulfilled (and if so, when—and how) is still an open question.

E. THE STATE OF THE ART

The high degrees of interest and promise in quantum communication and quantum computation are largely due to the central results of Bennett and Brassard [16], who designed a provably secure [123] communication protocol (BB84) using a quantum channel, and Shor [122], who devised an algorithm for finding the period of a sequence

exponentially faster than currently possible. This technique can be used to efficiently factor composite numbers or to calculate discrete logarithms, and so public-key cryptosystems and authentication protocols based on the supposed computational infeasibility of these number-theoretic problems, such as (among others) RSA [115], ElGamal [55], and the Digital Signature Algorithm (DSA)—the key element of the federal Digital Signature Standard (DSS) [59]—would therefore be rendered useless in the face of a quantum computational attack.

Various search algorithms proposed by Grover [67] and others raise the possibility of pattern-matching and recognition schemes of hitherto unimaginable power. Simulation of quantum mechanics and other physical systems [1], [2], [25], [135] could provide the tools necessary to design nanostructures [8]. A number of applications to statistical and numerical analysis (e.g., [3]), signal analysis, and so forth, have been discovered that are possible only in the realm of quantum computation. Exploitation of the universality and quasi-physical evolution properties of quantum cellular automata [39], [131] also holds theoretical and practical promise [14], [25], [103], [134]. Further significant theoretical advances are almost surely on the horizon.

Finally, although decoherence poses a formidable obstacle to the realization of quantum computers even with the use of quantum error-correcting codes, experimentalists have nevertheless recently constructed entangled states of four [116] and seven [85] qubits using ion trap [128] and liquid-state NMR [79] architectures, respectively.

II. BASIC THEORETICAL MODELS

A. LOGIC GATES

It so happens that *universal sets of logic gates* suffice to perform classical digital computation. If, for example, we consider the XOR (exclusive-OR or controlled NOT) and AND gates (which correspond, respectively, to addition and multiplication in the field $\mathbb{F}_2 \equiv \mathbb{Z}_2$), we can write any Boolean operation as an appropriate composition of these operations.

The principle of unitary evolution in quantum mechanics leads to time symmetry, however, and therefore any classical logic gates that we hope to carry over to the quantum regime must be reversible. Indeed, Landauer [82] and Bennett [15], in their analyses of fundamental lower bounds on heat dissipation resulting from computation, showed that models of classical reversible computers could be constructed. Key to such a construction is the augmentation of nominally one-output gates. For example, the augmented (reversible) XOR gate acts as $A, B \mapsto A, A \oplus B$. It is noteworthy that a quantum XOR gate was physically realized in an ion trap as long ago as 1995 [106]. This is by no means a trivial thing: a quantum XOR gate acting on the (normalized) state $|00\rangle + |10\rangle$ produces an EPR pair, and vice versa.

DiVincenzo showed [50] that two-qubit gates can be combined to form a universal three-qubit gate and hence that two-bit gates are universal for quantum computation. However, the gate decomposition provided therein was impractical. Sleator and Weinfurter proved that the gate with matrix in the canonical computation basis

$$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & e^{i\pi/4} \cos \pi\theta & e^{-i\pi/4} \sin \pi\theta \\ 0 & 0 & e^{-i\pi/4} \sin \pi\theta & e^{i\pi/4} \cos \pi\theta \end{pmatrix}$$

is universal [126]. Lloyd went further and demonstrated that almost any quantum gate with multiple inputs is universal. The key to this discovery was the realization that the algebra generated by two distinct n -qubit Hamiltonians is the space of Hermitian operators on \mathcal{H}^n , unless both Hamiltonians lie in a submanifold of positive codimension

[95]. Finally, it was demonstrated that one-bit and quantum XOR gates form a universal set [6].

Therefore, we can formally consider quantum computers as universal models for computation. The role in quantum computation analogous to that of Turing in the arena of classical computation could be said to have been filled by Feynman; it is to his construction we now turn.

B. FEYNMAN'S QUANTUM METAPROGRAM

The first step in devising algorithms and programs to run on a formal quantum computer was taken by Feynman [57], who constructed a metaprogram in the guise of a Hamiltonian on $n + k + 1$ qubits,

$$H = \sum_{i=0}^{k-1} a_{i+1}^* a_i A_{i+1} + a_i^* a_{i+1} A_{i+1}^* ,$$

where a_i^* , a_i are the creation and annihilation operators (sending $|0\rangle$ to $|1\rangle$ and $|1\rangle$ to $|0\rangle$, respectively) on the i th qubit, and A_i represents, for example, two-qubit gates acting on n qubits. In Feynman's proposal the $(k + 1)$ "program counter sites" (qubits) were initially set to $|0\rangle$ save for the initial qubit, which was set to $|1\rangle$; the Hamiltonian would then propagate this "cursor" state down the program counter sites, executing the A_i [98]. In this context quantum gates are easy to express: $a + a^*$ corresponds to NOT, $a^*a(b + b^*) + aa^*$ to XOR, and so forth. At the time, however, the utility of such a construct was unclear. This remained to be the case for nearly 10 years.

C. THE QUANTUM FOURIER TRANSFORM

Indeed, there is a vast gulf between models of universal computation or metaprograms and specific algorithms; despite an early recognition of the basic problem of how to generate constructive interferences in order to do anything useful with a quantum computer, no real progress in this area was made until the discovery (prompted by Shor's work) of a realistic quantum Fourier transform (QFT) by Coppersmith [44] and independently by Deutsch [47]. The QFT is a prototypical building block for quantum algorithms, and its central role can hardly be understated.

The quantum Fourier transform on n qubits is basically the Fourier transform on \mathbf{Z}_{2^n} :

$$QFT : \sum_{j=0}^{2^n-1} f(j) |j\rangle \mapsto \frac{1}{\sqrt{2^n}} \sum_{j=0}^{2^n-1} \sum_{k=0}^{2^n-1} e^{2\pi i j k / 2^n} f(j) |k\rangle = \frac{1}{\sqrt{2^n}} \sum_{k=0}^{2^n-1} \hat{f}(k) |k\rangle .$$

It can alternatively be described by the series composition of quantum gates

$$QFT = B \tilde{W}_1 S_{1,2} \dots S_{1,n-1} S_{1,n} \tilde{W}_2 \dots \tilde{W}_{n-2} S_{n-2,n-1} S_{n-2,n} \tilde{W}_{n-1} S_{n-1,n} \tilde{W}_n$$

where

$$B = \sum_{k'=0}^{2^n-1} |2^n-1-k'\rangle \langle k'|, \quad \tilde{W}_k = \bigotimes_{j=1}^{k-1} Id_j \otimes W \otimes \bigotimes_{j=k+1}^n Id_j,$$

$$S_{k,k'} |2^l\rangle = \bigotimes_{j=1}^{k'-1} Id_j \otimes \begin{pmatrix} 1 & 0 \\ 0 & e^{\frac{i\pi}{2^{k-k'}}} \end{pmatrix} \otimes \bigotimes_{j=k'+1}^n Id_j |2^l\rangle$$

(note that B is a bit reversal; this evokes the classical FFT [88]).

In practice (so to speak) the S -gates are simply discarded for small associated phases, resulting in a realizable QFT [44]. In fact, it turns out that the approximate QFT can actually improve performance for periodicity estimates in the presence of decoherence [7]. Finally, it is worth noting that the QFT and Walsh-Hadamard transform act identically on $|0\rangle$.

Various generalizations of the QFT have been outlined. For example, a quantum wavelet transform has been developed [52]. Kitaev [83] constructed an analog of the QFT for finite Abelian groups (actually cyclic groups \mathbb{Z}_p of prime order; however, by the fundamental theorem of abelian groups [54], this is sufficient). A quantum network of gates could be designed to perform such a transform efficiently along lines not entirely dissimilar to the Coppersmith construction. In general, a QFT runs exponentially faster than a classical FFT—indeed, the QFT requires only a quadratic number of gates; this improvement over the $O(n2^n)$ operations required for the equivalent FFT is a central result in quantum complexity theory. Tighter bounds on the circuit complexity of the QFT can be found in [42].

D. QUANTUM ALGORITHMS FOR SPECIAL ORACLE PROBLEMS

It has long been recognized that the augmentation of a classical Turing machine with an *oracle* capable of addressing queries with respect to nonrecursive functions (i.e., functions not specified by a formula or algorithm but rather as a “black box”) would allow the efficient solution of problems beyond the scope of an ordinary classical Turing machine [11]. With this in mind, Bernstein and Vazirani [22], Deutsch and Jozsa [46], and Simon [125] exploited quantum parallelism to exhaust an oracle and thereby arrive at quantum algorithms with better performance than is classically possible. To illustrate the

nature of the oracle problem in quantum computing we sketch the Deutsch-Josza algorithm (DJ).

The context of DJ is specified by a nonrecursive function $f: \mathbf{Z}_{2^n} \rightarrow \mathbf{Z}_2$ which is promised or assumed a priori to be either the zero function or to take each of the values 0 and 1 2^{n-1} times (in which case we refer to it as *balanced*). DJ differentiates between the two cases as follows:

- Step 0: Initialize an $n + 1$ qubit string $|0\dots 0\rangle|1\rangle$ (we could also write this as $|0\rangle|1\rangle$).
- Step 1: Apply the Walsh-Hadamard transform to each register to get

$$\frac{1}{\sqrt{2^n}} \sum_{i=0}^{2^n-1} |i\rangle \otimes \frac{1}{\sqrt{2}} (|0\rangle - |1\rangle) .$$

- Step 2: Apply the function via $|i\rangle|j\rangle \mapsto |i\rangle|j \oplus f(i)\rangle$ to get

$$\frac{1}{\sqrt{2^n}} \sum_{i=0}^{2^n-1} (-1)^{f(i)} |i\rangle \otimes \frac{1}{\sqrt{2}} (|0\rangle - |1\rangle) ,$$

since

$$\frac{1}{\sqrt{2^n}} \sum_{i=0}^{2^n-1} (-1)^{f(i)} |i\rangle \otimes \frac{1}{\sqrt{2}} (|0 \oplus f(i)\rangle - |1 \oplus f(i)\rangle) = \frac{1}{\sqrt{2^n}} \sum_{i=0}^{2^n-1} |i\rangle \otimes \frac{1}{\sqrt{2}} (|0\rangle - |1\rangle) .$$

- Step 3: Invert the Walsh-Hadamard transform on the first register to get

$$|dj\rangle \otimes \frac{1}{\sqrt{2}} (|0\rangle - |1\rangle) ,$$

$$\text{where } |dj\rangle = \begin{cases} |0\rangle & \text{if } f \text{ is zero} \\ |b\rangle \neq |0\rangle & \text{if } f \text{ is balanced} \end{cases} .$$

- Step 4: Measure the first register.

Because of our assumption that the function is either zero or balanced, we can determine with probability 1 the answer to the question of which type of function we actually have after performing the DJ algorithm. This is somewhat unusual and is an artifact of the “promise” made in the problem. It is noteworthy that a classical solution to the problem requires $O(2^{n-1})$ steps.

Simon’s algorithm is in the same spirit as DJ but is also slightly more subtle. For a nonrecursive function $f: \mathbf{Z}_{2^n} \rightarrow \mathbf{Z}_{2^n}$ which is assumed a priori to be one of the two cases, Simon’s algorithm determines (by using the QFT) whether f is one-to-one or two-to-one [4], [125]. Despite the exponential speedups these algorithms offer, however, they are

essentially toy models; the recent surge in interest in quantum computation derives from a far more useful quantum algorithm.

III. FACTORING ON A QUANTUM COMPUTER

A. SHOR'S ALGORITHM

Factoring is hard and important; we devote Appendix A to the explanation and ancillary results.

Shor [122] devised an ingenious method for factoring based upon two principles: one a known number-theoretical technique [104], the second quantum-computational. We outline his results correspondingly.

If we have a number N (which we assume not to be a prime power) and the order $r(x)$ (i.e., the smallest integer $r(x)$ such that $x^{r(x)} \equiv 1 \pmod{N}$) of any element x in the multiplicative group

$$\mathbf{Z}_N^* = \{a \in \mathbf{Z}_N \mid \gcd(a, N) = 1\} ,$$

we can consider $\gcd(x^{r(x)/2} - 1, N)$ for x random and such that $x^{r(x)/2} \pmod{N} \neq N - 1$ and $r(x) \equiv 0 \pmod{2}$. In this event it follows that since $(x^{r(x)/2} - 1)(x^{r(x)/2} + 1) \equiv 0 \pmod{N}$ we obtain a nontrivial factor. Shor's algorithm determines the order $r(x)$ as follows:

- Step 0: Initialize a $2n$ (where $N < 2^n$) qubit string $|0, 0\rangle$.
- Step 1: Apply the QFT to the first register to get

$$\frac{1}{\sqrt{2^n}} \sum_{a=0}^{2^n-1} |a, 0\rangle .$$

- Step 2: Compute $x^a \pmod{N}$ by using quantum gates that efficiently perform binary modular exponentiation for fixed x, N built into the gate structure—which can, in turn, be efficiently constructed from gates performing binary modular addition (such gates are described in [10], [53], [132]) for each element of the superposition:

$$\frac{1}{\sqrt{2^n}} \sum_{a=0}^{2^n-1} |a, x^a \pmod{N}\rangle .$$

- Step 3: Apply the QFT on the first register to get

$$|\Psi\rangle = \frac{1}{2^n} \sum_{a=0}^{2^n-1} \sum_{c=0}^{2^n-1} e^{2\pi i ac/2^n} |c, x^a \pmod{N}\rangle .$$

- Step 4: Measure both registers.

The probability of measuring the state $|c, x^{a'} \bmod N\rangle$ is given by

$$\begin{aligned} & \text{Tr}(|c, x^{a'} \bmod N\rangle\langle c, x^{a'} \bmod N| \Psi\rangle\langle\Psi|) \\ &= \frac{1}{2^{2n}} \left| \sum_{a=br(x)+a'}^{2^n-1} e^{2\pi i ac/2^n} \right|^2 = \frac{1}{2^{2n}} \left| \sum_{b=0}^{\lfloor 2^n-1-a'/r(x) \rfloor} e^{2\pi i b\{r(x)c\}/2^n} \right|^2, \end{aligned}$$

where $\{r(x)c\} \equiv rc \bmod 2^n$, $-2^{n-1} < \{r(x)c\} \leq 2^{n-1}$. If $|\{r(x)c\}| \leq r(x)/2$, it turns out that this probability is approximately

$$\left| \frac{1}{r(x)} \int_0^1 e^{2\pi i u\{r(x)c\}/r(x)} du \right|^2 \geq \left(\frac{1}{\pi\{r(x)c\}} \right)^2 \geq \left(\frac{2}{\pi r(x)} \right)^2,$$

where we have neglected lower order terms. In the limit of large N (and hence n), the probability distribution becomes a *Dirac comb* with spikes at values of c where there exists d such that $c = \lfloor d2^n / r(x) \rfloor$. It follows that [since as can be shown there are $r\phi(r)$ spikes] the probability of measuring a spike state is asymptotically [70]

$$\frac{4r\phi(r)}{r^2\pi^2} \geq \frac{\delta}{\pi^2 \log \log r}$$

for a constant δ . Therefore, in principle the measurement problem is solved at this point, and from a measurement of a spike we can determine the order $r(x)$ using techniques of continued fractions [70], [86].

Taking into account repeated trials, Shor's algorithm requires $O((\log N)^2 \log \log N \log \log \log N) \equiv O(n^2 \log n \log \log n)$ steps; additional polynomial post-processing time is necessary to efficiently determine a factor from the order of a suitable element classically. The majority of the quantum processing time is spent in performing modular exponentiation; more efficient techniques for this can further enhance Shor's algorithm.

B. FACTORING AS AN INSTANCE OF THE ABELIAN STABILIZER PROBLEM

Kitaev [83] generalized the factoring and discrete logarithm (Appendix A) problems in the context of the *abelian stabilizer problem (ASP)*: given an action α of \mathbf{Z}^k on $M \subset \mathbf{Z}_2^n$, that is, given $\alpha: \mathbf{Z}^k \times M \rightarrow M$ with $\alpha_{g+h}(a) = \alpha_g \alpha_h(a)$, determine a basis of the stabilizer $St_\alpha(a) \equiv \{g: \alpha_g(a) = a\}$. The problem is well posed since the stabilizer is a finite-rank subgroup of \mathbf{Z}^k [54]. To see that factoring is an instance of the ASP, consider

$M = \mathbf{Z}_N$, $G = \mathbf{Z}_N^*$ and an action defined by $\alpha_m^x(a) = x^m a$. In this context a basis of the stabilizer $St_\alpha(1)$ gives the order $r(x)$.

We present a sketch of the quantum AS algorithm for factoring. Consider first the quotient group $E \equiv \mathbf{Z}/St_\alpha(1) \equiv \mathbf{Z}_{r(x)}$ and its *character group* [100] \hat{E} of homomorphisms from E to the circle. A character χ_h is now characterized by a rational number $h/r(x)$ between 0 and 1 [i.e., to specify a homomorphism χ_h from a cyclic group, which we can represent as roots of unity, to the circle we need only the number (h) of times χ_h wraps around the circle]. Indeed, a cyclic group is isomorphic to its character group [100] and so if we can determine the wrapping number, h , of a generator, then the factoring ASP is effectively solved.

Toward this end we can consider elements of E as shift operators on the orbit of 1 given by $N \equiv \{\alpha_m^x(1) : m \in \mathbf{Z}\} = \{x^m\}$; the solution of (the factoring instance of) the ASP depends on measuring an eigenvalue of such a shift. Kitaev's scheme uses for unitary shift operators on $\mathbf{Z}_{r(x)}$ with eigenvectors $|\psi_{h,r(x)}\rangle$ and corresponding eigenvalues $e^{-2\pi i h/r(x)}$ a transformation such as

$$(Id \otimes W)^{-1} \left(Id \otimes \begin{pmatrix} 1 & 0 \\ 0 & e^{-2\pi i h/r(x)} \end{pmatrix} \right) (Id \otimes W) ,$$

which behaves as

$$|\psi_{h,r(x)}\rangle \otimes |0\rangle \mapsto |\psi_{h,r(x)}\rangle \otimes e^{-\pi i h/r(x)} (\cosh(-\pi i h/r(x))|0\rangle + \sinh(-\pi i h/r(x))|1\rangle) ,$$

to bias the control (second) register. By measuring enough of these identically prepared states for binary powers of a shift operator (and performing some subtle handwaving), an observer can approximate the phase of an eigenvalue to any desired accuracy with high probability in polynomial time [4], [83]. (This methodology is also employed for Kitaev's QFT.) This is an instance of so-called *eigenvalue estimation*, which also appears in algorithms for quantum mechanical simulation (see Section V.B.), for example. The final step in Kitaev's algorithm as we present it now is to prepare a uniform superposition of all the shift eigenvectors (which can also be done efficiently) and use the biasing scheme to measure a given value h .

IV. SEARCH ALGORITHMS

A. FINDING A NEEDLE IN A QUANTUM HAYSTACK

It is intuitively obvious why a classical search routine applied to an unstructured list of N objects must take at least $O(N)$ steps: if our list has no internal structure then we must perform an exhaustive search. (If, on the other hand, we can progressively subdivide our list, for example in a balanced binary tree, then we can perform a classical search in $O(\log N)$ steps.)

Grover discovered the amazing result that a *quantum* search for such a “needle in a haystack” could be performed in $O(\sqrt{N})$ steps [67] (in fact, it has been shown that this is also a lower bound [28], [136]). While perhaps less intriguing on the surface than Shor’s factorization algorithm, Grover’s haystack search is certainly more versatile; the underlying technique of *amplitude amplification* can be brought to bear on a host of problems.

Grover’s haystack search proceeds as follows. We are given an oracle $f : \mathbb{Z}_2^n \rightarrow \mathbb{Z}_2$ with *only* one target state t such that $f(t) = 1$. Consider the states

$$|a\rangle = \frac{1}{\sqrt{2^n}} \sum_{i=0}^{2^n-1} |i\rangle, \quad |u\rangle = \frac{1}{\sqrt{2^n-1}} \sum_{i \neq t} |i\rangle = \frac{1}{\sqrt{2^n-1}} \sum_{i=0}^{2^n-1} (1 \oplus f(i)) |i\rangle$$

and the operator R corresponding to a rotation by $\theta \equiv \cos^{-1} \langle a|u \rangle$ on the two-dimensional subspace T spanning $|a\rangle$ and $|u\rangle$ (equivalently, $|a\rangle$ and $|t\rangle$). Furthermore, denote by $M_{|b\rangle}$ the reflection $Id - 2|b\rangle\langle b|$ about the subspace spanned by a single state $|b\rangle$. Though we have no direct knowledge of T , it is clear by inspection that $R^2 = M_{|b\rangle} M_{|u\rangle}$. Moreover, $WM_{|0\rangle}W = W^2 - 2|a\rangle\langle a| = Id - 2|a\rangle\langle a| = M_{|a\rangle}$ and $M_{|u\rangle}|i\rangle = (-1)^{f(i)} |i\rangle$. Now $\cos \theta = \langle a|u \rangle = \sqrt{(2^n - 1)/2^n}$, so $\theta \approx \sin \theta = 2^{-n/2}$, and it follows that

$$\left(M_{|a\rangle} M_{|u\rangle} \right)^{\lfloor \pi \sqrt{2^n} / 4 \rfloor} |a\rangle \approx |t\rangle.$$

Summarizing, we can perform a quantum search as follows:

- Step 0: Initialize a n qubit string $|0\rangle$.

- Step 1: Apply the Walsh-Hadamard transform to get

$$|a\rangle = \frac{1}{\sqrt{2^n}} \sum_{i=0}^{2^n-1} |i\rangle .$$

- Step 2: Rotate to get

$$\left(M_{|a\rangle} M_{|u\rangle} \right)^{\lfloor \pi\sqrt{2^n}/4 \rfloor} |a\rangle \approx |t\rangle .$$

- Step 3: Perform a measurement.

This yields the correct state with high probability (asymptotically 1).

The repeated application of a small rotation is generally referred to as amplitude amplification. It is a powerful and versatile method. *It could be said at present that fast quantum algorithms invoke an oracle, a QFT, amplitude amplification, or eigenvalue estimation; as it stands, these are the tools of the trade in quantum computing.* Below we provide a sketch of the generalized Grover search algorithm that serves to illustrate amplitude amplification in a more general format.

B. GENERALIZED UNSTRUCTURED AND UNSTRUCTURED PARALLEL QUANTUM SEARCHING

Grover's quantum search generalizes to multiple target states and arbitrary operators and initial superpositions, as shown by Gingrich, Williams, and Cerf [62]. Therefore, quantum search is a viable subroutine for more general programs. Following [62], we outline the algorithm:

- Step 0: Initialize a qubit string

$$2^{-1/2}(|a\rangle|0\rangle - |a\rangle|1\rangle) = 2^{-1/2}((|a\rangle - |t\rangle)|0\rangle - (|a\rangle - |t\rangle)|1\rangle + |t\rangle|0\rangle - |t\rangle|1\rangle) ,$$

where $|t\rangle$ is the superposition of target states.

- Step 1: Apply the oracle to get

$$2^{-1/2}((|a\rangle - |t\rangle)|0\rangle - (|a\rangle - |t\rangle)|1\rangle - |t\rangle|0\rangle + |t\rangle|1\rangle) = 2^{-1/2}(|a\rangle|0\rangle - |a\rangle|1\rangle - 2(|t\rangle|0\rangle - |t\rangle|1\rangle)) ,$$

which is equivalent to applying the inversion operator $Id - 2|t\rangle\langle t|$. (Though equivalent, this is not the same thing: we do not know what the target states are.)

- Step 2: Pick an inversion state $|b\rangle$ and apply the operator $-(Id - 2|b\rangle\langle b|)$ an appropriate number of times (depending on both the initial and inversion states).

- Step 3: Repeat Steps 1–2; after an appropriate number of iterations, perform a measurement.

In this general setting there is a phase condition that governs the probability of success for amplitude amplification; this can be used to construct appropriate inversion states [76].

Gingrich, Williams, and Cerf also analyzed punctuation (premature halting after a submaximal number of rotations or inversions) and parallelization of the generalized Grover search routine. They found first that punctuation actually speeds up the routine and is maximized (12-percent fewer rotations/inversions than the Grover algorithm, which, it should be recalled, is therefore 12-percent faster than the quickest possible complete quantum search) when the probability of successful search is 84 percent, and second, that parallelization, even in the optimal case, is useful primarily as a stay against decoherence; indeed, since the gain in time turns out to be $O(\sqrt{k})$ for a parallel quantum search by k quantum computers, the cumulative time of the parallel search actually exceeds that of a single-agent search by a factor $O(\sqrt{k})$.

C. STRUCTURED QUANTUM SEARCHING

Hogg [74] noted that the potential of quantum computation is difficult to evaluate on the basis of overly specific (e.g., factoring) or general (e.g., unstructured search) algorithms. An intermediate problem in this context is the so-called *random K-SAT (satisfiability) problem* [81], in which a solution to a formula

$$F = \bigwedge_{i=1}^m C_i = \bigwedge_{i=1}^m \bigvee_{j(i)=1}^K b_{ij(i)}$$

satisfies all m clauses of logical-ORs and NOTs of K (of n total) Boolean variables. Here, the b -terms denote *literals* (i.e., either a variable or its negation, with equal probabilities). Varying the ratio of clauses to total variables leads to a phase transition in the problem difficulty [82], [105]: if this ratio is small, then many solutions exist and the problem is easy; if it is large, then no solution exists and the decision problem [11] is easy. For intermediate values, however, the problem is difficult. Hogg studied quantum-computational approaches to K -SAT for the maximally constrained case (i.e., for the largest number of clauses possible in order to retain a solution) [73] and near the phase transition [74].

(The random K -SAT problem has recently been extensively investigated. In particular, composite problems interpolating between $K = 2$ [which has a linear time

solution] and $K = 3$ [which is NP-complete] and the accompanying phase transition [82], [105] have been the source of considerable interest owing to their status as transitional prototype problems.)

The general method of attack for such a problem is to perform a quantum search of partial solutions and use this to restrict the remainder of the search space [36]. Restricting the search space improves classical and quantum algorithms by raising the execution time of both to the same power $\alpha < 1$; the quadratic speedup afforded by Grover's algorithm for haystack searches holds for structured searches. (Search algorithms which run in linear time classically—such as maximally constrained K -SAT—experience a speedup to constant time [73].) Successive restrictions do the same, with progressively smaller exponents. In this sense the structured search can be seen as a generalized dynamic tree search, with the limiting case of a fixed (e.g., binary) tree search requiring logarithmic time (which is consistent with the power-law scaling). With this in mind, the quadratic speedup may be said to be a universal feature of quantum versus classical structured search methods [29], [36].

It is interesting that spin glass models have been applied to the study of random K -SAT [24]. A vast undiscovered country with the potential to provide feasible solutions to general hard problems—as well as fundamental information on the difficulties and limits of computation—lies at the confluence of spin glasses, satisfiability and other hard problems, and quantum computation. In particular, analyses of heuristic searches and optimization routines hold promise for evaluating quantum computation [74], [75]. Such heuristics can also find application in (e.g.) the traveling salesman problem [75] and game-theoretic routines (e.g., minimax or checkmate searches); the possibility of performing wargaming or logistical computations on quantum computers merits examination. Even a protocol for appointment scheduling exists [32].

V. OTHER SELECTED ALGORITHMS

A. INTEGRATION

Abrams and Williams [3] have refined a novel technique of Grover [68] for approximating integrals (equivalently, the mean of a sequence) iteratively on a quantum computer. Their approach is not dissimilar to the Monte Carlo method of integration (where the integral is approximated by the function values at random points) but offers a quadratic speedup over it (and a speedup over deterministic methods that is exponential in the dimension of the function space); indeed, their method is based on the same amplitude amplification principle as a quantum search algorithm.

Let E be an estimated value of $S = \langle f \rangle$, the average of a step function (without loss of generality assumed to have range contained in the unit interval over a uniformly subdivided unit d -cube):

$$S = \frac{1}{M^d} \sum_{a_1, a_2, \dots, a_d=1}^M f(a_1/M, a_2/M, \dots, a_d/M) = \frac{1}{M^d} \sum_{a_1, a_2, \dots, a_d=1}^M \tilde{f}(a_1, a_2, \dots, a_d) .$$

Put $D = S - E$ and $g = \tilde{f} - E$, so that $D = \langle g \rangle$.

Use the Walsh-Hadamard operator to prepare the state

$$\frac{1}{\sqrt{M^d}} \sum_{a_1, a_2, \dots, a_d=1}^M |a_1, a_2, \dots, a_d\rangle |0\rangle ,$$

and apply a rotation to the second register to get

$$\frac{1}{\sqrt{M^d}} \sum_{a_1, a_2, \dots, a_d=1}^M \sqrt{1 - g^2(a_1, a_2, \dots, a_d)} |a_1, a_2, \dots, a_d\rangle |0\rangle + g(a_1, a_2, \dots, a_d) |a_1, a_2, \dots, a_d\rangle |1\rangle ,$$

and apply the inverse Walsh-Hadamard operator. It so happens that the probability of observing $|0\rangle|1\rangle$ is D . By performing the above procedure from scratch enough times we can therefore determine D (and hence E and thence S). If, however, we consider the rotation of the first qubit as the rotation in a quantum search algorithm and $|0\rangle|1\rangle$ as our target state, we can perform an amplitude amplification by repeating this sequence of operations. This is quadratically faster than sampling—of order the inverse (as opposed to the inverse square) of the desired accuracy.

Alternatively, *quantum counting* [29], a variant of amplitude amplification, can be used. In this scenario an auxiliary parameter q with integer values from 1 to Q (determined by the desired degree of accuracy) is introduced so that the mean of the Boolean function

$$b(a_1, a_2, \dots, a_d, q) \equiv \begin{cases} 1 & \text{if } q \leq Q \cdot \tilde{f}(a_1, a_2, \dots, a_d) \\ 0 & \text{if } q > Q \cdot \tilde{f}(a_1, a_2, \dots, a_d) \end{cases}$$

equals S . The number r of solutions of $b = 1$ can be counted by invoking amplitude amplification because the amplitude of a single rotation turns out to be proportional to the square root of r . If a superposition of states corresponding to successive powers of the basic amplifying rotation is created, then r can be determined by performing a QFT on a register indicating the power or number of rotations. (See V.B for a similar procedure in the context of eigenvalue estimation as applied to quantum mechanical simulation.) The performance of this algorithm also scales inversely to the desired accuracy. Details can be found in [3].

As it turns out, the principal difficulty with Monte Carlo integration and classical implementations of probabilistic algorithms in general is generating truly random points: as Knuth points out [88], generating even suitably good pseudorandom values is very difficult (it could be said that all the cryptographers in the world have failed to devise a pseudorandom number generator which performs well enough to satisfy themselves). It is also interesting that the runtime of the Monte Carlo method and general amplitude amplification algorithms depend not on the size of the problem per se but rather on the desired accuracy (a function of the number of measurements required for a sufficiently high probability of getting the correct answer, which *will* in some sense depend on the problem size).

B. SIMULATION OF LOCAL QUANTUM SYSTEMS

Quite possibly the most important—and most immediately realizable—application of quantum computing is the purpose which Feynman originally envisioned: simulating quantum mechanics [96]. Lloyd [95] provided the basic theoretical framework for directly simulating *any* local quantum system (such as an Ising [63], [93] or lattice gauge [110] model) with an *exponential* improvement over classical simulation. Interestingly, the converse is commonplace: quantum computers are frequently simulated via Ising spin models.

Lloyd's basic setup is as follows. A local quantum system with Hamiltonian

$$H = \sum_{j=1}^l H_j ,$$

such that each term in the sum acts on a local Hilbert space of finite dimension $d_j = \dim(\text{supp}(Id - H_j))$ and has a time-evolution operator which is decomposed over short time intervals via the Campbell-Baker-Hausdorff formula [117]:

$$e^{iHt} = \left(\prod_{j=1}^l e^{iH_j \Delta\tau} \right)^{t/\Delta\tau} + \sum_{j>k} [H_j, H_k] \frac{t\Delta\tau}{2} + O\left(\frac{te^{\|H\|\Delta\tau}}{\Delta\tau}\right),$$

where simulating each operator $e^{iH_j \Delta\tau}$ requires $O(d_j^2)$ operations. Under this decomposition the total number of operations for the time evolution operator is $O(tl \max(d_j^2)/\Delta\tau)$; with this in mind, we require that the number l of terms in the Hamiltonian should scale as a polynomial function of the number of variables or particles. Moreover, we can use this complexity analysis to specify the number (n) of time slices required for a simulation of a given accuracy. Finally, this formalized scheme can accommodate environmental interactions either by including extra terms in the model Hamiltonian or, more elegantly, by simply scaling the computer's environment suitably and exploiting, say, decoherence in the physical system to simulate decoherence in the model system.

In [1] Abrams and Lloyd also outlined an efficient polynomial algorithm for producing an antisymmetrized superposition of states representing the initial state of a fermionic system of k particles, which can then be time-evolved as above. As an alternative, they propose a model based on a quantum field-theoretic or second quantized [110] formalism which is sometimes in principle (i.e., when the number of particles $k \ll m$, where m is the number of single particle states) more efficient, owing to the Pauli exclusion principle and the concomitant encoding of the state of the fermionic system into a bit vector of length m . The field-theoretic formulation and its corresponding time evolution are more involved, however; refer to [1] for the details.

Abrams and Lloyd, following Cleve et al. [41], also proposed a more explicit methodology for using the QFT to find eigenvalues and infer eigenvectors of evolution operators (hence also of Hamiltonians) in polynomial time [2]. We sketch the procedure for determining the eigenvalues.

- Step 0: Initialize a qubit string $|0\rangle|\psi\rangle$ of length $m + l$, where $|\psi\rangle$ is an approximate eigenvector of the time evolution operator $U = e^{-iHt}$:

$$|\langle \phi_k | \psi \rangle|^{-2} = O(\text{poly}(l)) ,$$

where ϕ_k , $\lambda_k = e^{i\omega_k}$ are the eigenvectors and eigenvalues of U . (Such an approximate eigenvector can be generated in polynomial time by invoking a classical approximation.)

- Step 1: Apply the QFT on the index register, obtaining the state

$$\frac{1}{\sqrt{2^m}} \sum_{j=0}^{2^m-1} |j\rangle |\psi\rangle .$$

- Step 2: Produce the state

$$\frac{1}{\sqrt{2^m}} \sum_{j=0}^{2^m-1} |j\rangle U^j |\psi\rangle = \frac{1}{\sqrt{2^m}} \sum_{j=0}^{2^m-1} |j\rangle U^j \sum_k \langle \phi_k | \psi \rangle |\phi_k\rangle = \frac{1}{\sqrt{2^m}} \sum_k \langle \phi_k | \psi \rangle \sum_{j=0}^{2^m-1} |j\rangle e^{i\omega_k j} |\phi_k\rangle$$

by, for example, binary exponentiation of the evolution operator conditioned on the first register. (This step is reminiscent of quantum counting, cf. IV.A.)

- Step 3: Apply the QFT on the index register again to get

$$\frac{1}{2^m} \sum_k \langle \phi_k | \psi \rangle \sum_{j,n=0}^{2^m-1} e^{2\pi i j n / 2^m} e^{i\omega_k j} |n\rangle |\phi_k\rangle = \frac{1}{2^m} \sum_k \langle \phi_k | \psi \rangle \sum_{j,n=0}^{2^m-1} e^{i(\omega_k + 2\pi n / 2^m) j} |n\rangle |\phi_k\rangle .$$

- Step 4: Perform a measurement on the index register. The probability of measuring $|\phi_k\rangle$ is

$$|\langle \phi_k | \psi \rangle|^2 .$$

A polynomial number of repetitions then gives the eigenvalues satisfying the approximation requirement in Step 0 with an accuracy which scales as 2^{-m} . By performing the CBH decomposition as above, the evolution operator can be realized and the problem solved in polynomial time.

Kitaev's algorithm [83] is in the same spirit as the Abrams-Lloyd algorithm: both are instances of the eigenvalue estimation meta-algorithm. The latter, however, can be exploited to determine physical quantities which are also functions of the eigenvectors, such as charge density and momentum distributions or correlation functions [2].

C. QUANTUM CELLULAR AUTOMATA

It has long been known that cellular automata (CA)—and particularly reversible CA—are universal models of computation [61], [131]. As a consequence, certain so-called *lattice-gas cellular automata* (LGCA) and related ballistic systems can perform computation as a manifestation of their quasi-physical dynamics [39]. Therefore, it is natural in this context to wonder whether a direct physical incarnation of a CA (or LGCA) can be realized.

It so happens that quantum cellular automata (QCA) and quantum lattice gas automata (QLGA) can be defined [103]. The collision operator or transition rule is given as a sort of S-matrix which acts on superpositions of incoming states and yields superpositions of outgoing states at each time step, rather than a deterministic or probabilistic rule operating on fixed states which cannot be superposed.

LGCA such as the FHP hydrodynamical models and the lattice-Boltzmann models obtained by averaging the quantities in the collision operator thereof [39] can simulate instances of the Navier-Stokes equations. In much the same way, QLGA can simulate the N -body Schrödinger equation [24] or the N -body Dirac equation in one dimension [58], [103]. (Interestingly enough, QCA have even been considered as a vehicle to simulate Navier-Stokes [135].)

The details of simulating the N -body Schrödinger equation in d dimensions are involved; for them refer to [25]. The generic dynamics are given by equations of the form

$$\psi_{i_1 \dots i_N}(x_1 + \varepsilon_{c_1}, \dots, x_N + \varepsilon_{c_N}, t+1) = \prod_{k,l} S_{i_k j_l} \psi_{j_1 \dots j_N}(x_1, \dots, x_N, t) ,$$

where ε_{c_i} is a lattice vector. (A general LGA can also be put in this form; the key is that the collision operator here is an S-matrix and not a classical “billiard-ball” collision operator such as arises in, e.g., FHP models.) It is worth noting that this model allows for the inclusion of a general potential (via multiplication of the S-matrix by a position-dependent phase) and for hard-Bose or Fermi statistics.

Interest in QCA is not just related to physical simulation, however. The universality properties of automata suggest architectures for actual quantum computers in much the same way that the universal Turing machine might have suggested a real computer using a magnetic tape. Benjamin and Johnson [14] have developed a prototype scheme for quantum computation that exploits conventional as well as quantum parallelism by considering various types of qubits or “cells.” A certain cell type is associated with a distinct energy gap between its two states; the gate architecture in the prototype is given by the spatial configurations of various cell types into networks. However, the geometry of the network per se is inessential; the key is where the various cell types are located within the geometry.

If, though, we can consider multiple-state quantum systems (“qubytes”) such that we can restrict the allowable states selectively and independently from qubyte to qubyte in a reasonable manner, we can avoid building a specialized network for any one particular algorithm. (This generalization is analogous to the relationship between, say, a

mechanical differential equation solver and a PC loaded with a differential equation software package.)

A specific example of such a cellular quantum computer utilizes two types of cells and (essentially) six local updates; this setup is sufficiently general to provide a universal quantum computer which can be massively parallelized spatially (“pipelined”). By performing massively parallel independent amplitude amplifications and sequentially measuring the network output nodes as successive amplifications take place on the remainder, the runtime of a search would reduce in line with the later results of [62] for optimal parallel search.

VI. QUANTUM INFORMATION THEORY

A. QUANTUM COMMUNICATION CHANNELS

Consider for the moment a classical (n, k) linear binary code C (Appendix B). In this context, error correction and decoding are deterministic processes. That is, a given k -block has a unique associated codeword (i.e., a canonical representative of the code co-set of which the block is a member) which is then transmitted, received, and decoded uniquely. If the transmitted and received n -blocks belong to the same co-set of the equivalence relation induced by C , then the error-correction mechanism will succeed.

In the quantum regime, we can consider the coding problem in an ensemble as well as the possibility of decoding errors (note that encoding errors and transmission errors are effectively the same thing). Now the channel can be represented by the *S-sequence*:

$$S: \mathcal{H}^k \otimes \mathcal{H}^{n-k} \xrightarrow{\mathcal{E}} \mathcal{H}^n \xrightarrow{\tau} \mathcal{H}^n \xrightarrow{\mathcal{D}} \mathcal{H}^k \otimes \mathcal{H}^{n-k} \xrightarrow{\delta} \mathcal{H}^n \xrightarrow{\mathcal{D}} \mathcal{H}^k \otimes \mathcal{H}^{n-k},$$

where the last two stages of the sequence may repeat several times if detectable errors occur during the decoding process. The operators are unitary and act on the n -qubit Hilbert space, and $\mathcal{D} = \mathcal{E}^{-1}$. If we assume, for example, that decoding is error free and only one-qubit rotation errors (distributed symmetrically over, say, a parameter space $-1 < \theta < 1$) occur, the *S-sequence* becomes

$$S = \mathcal{D} \circ \bigotimes_{i=1}^n R_{\theta_i} \circ \mathcal{E}.$$

(The tacit assumption here is that most of the terms in the tensor product are effectively identity operators.)

Suppose further that the (classical) probability of a single one-qubit error occurring is p (presumably a time-dependent function) and that errors are independent. Then the probability of m one-qubit errors is given by a binomial distribution, and we find that the ensemble output of the channel can be represented by an ensemble density matrix:

$$\sum_{a=0}^{2^k-1} \Pr(|a, 0\rangle) \rho_{|a, 0\rangle, \theta} = \sum_{a=0}^{2^k-1} \Pr(|a, 0\rangle) \sum_{m=0}^n \sum_{\{j\}_m} p^m (1-p)^{n-m} \left| b(a)_{\{j\}_m} \right\rangle \left\langle b(a)_{\{j\}_m} \right|,$$

$$\left| b(a)_{\{j\}_m} \right\rangle = S_{\{j\}_m} |a, 0\rangle \equiv \mathcal{D} \circ \bigotimes_{i=1}^m R_{\theta_{j_i}} \circ \mathcal{E} |a, 0\rangle; \{j\}_m \equiv \{j_1 < \dots < j_m\},$$

(where the strike through the tensor product indicates suppression of factors that are effectively identity operators). The *von Neumann entropy* is defined for a density matrix ρ as $H(\rho) = -\text{Tr}(\rho \log_2 \rho)$ (which is well defined since the density matrix is a positive operator and the techniques of spectral functional calculus can be brought to bear on it [43]). It turns out [118] that the von Neumann entropy is the appropriate generalization of the classical Shannon entropy insofar as it is an ideal lower bound on the expected length of a qubit string encoding the ensemble described by the density matrix.

The *Levitin-Holevo upper bound* (see, e.g., [71]) on the classical mutual information is given by

$$H\left(\sum_a \Pr(a) \rho_a\right) - \sum_a \Pr(a) H(\rho_a);$$

if there are no transmission errors in our example, then the density matrices represent pure states, the second term vanishes, and what remains is just the von Neumann entropy. (However, the above expression holds even when the input states are mixed.) It was shown in [71] that by coding properly (and in the absence of noise) the classical mutual information per qubit can be brought arbitrarily close to the ensemble von Neumann entropy. From this it follows that the natural definition for a *quantum channel capacity* is the maximum possible von Neumann entropy (since this is also the maximum classical mutual information).

The *fidelity* of our ensemble S -sequence in our example is now

$$F \equiv \left\langle \text{Tr}(\langle a, 0 | S | a, 0 \rangle \rho) \right\rangle$$

$$= \frac{1}{2^k} \sum_{a=0}^{2^k-1} \Pr(|a, 0\rangle) \int_{-1}^1 \text{Tr} \sum_{m=0}^n \sum_{\{j\}_m} p^m (1-p)^{n-m} \left| b(a)_{\{j\}_m} \right\rangle \left\langle b(a)_{\{j\}_m} \right| \left\langle a, 0 \right| b(a)_{\{j\}_m} \rangle p_\theta d\theta.$$

This and like expressions are realistic gauges of a quantum error-correcting code or a quantum channel [84], and its calculation requires for a given code and signal ensemble only the probability of a single error occurring and a probabilistic description of a one-qubit error operator, both of which can be determined through experiment. Moreover, it is clear that a similar (but more complicated) expression holds when allowing errors in the encoding and decoding stages of the S -sequence or many-qubit errors. These extensions are straightforward. Not so straightforward is the inclusion of decoherence as a

transmission error, which is of primary concern because the process of decoherence is effectively nonunitary and in particular can result in what amounts to a non-invertible S -sequence in which the input signal space is collapsed onto a subspace. In any case, our definition of fidelity is essentially the average fidelity of [34]; similar definitions can be found throughout the literature, almost all with particular variations.

Schumacher showed in [118] that a quantum analog of Shannon's noiseless coding theorem holds. In particular, let a quantum channel have the (quantum) capacity C and a quantum source the entropy per unit time H . If $H \leq C$, there exists a coding system such that the output of the source can be transmitted (but not copied) over the channel with a fidelity arbitrarily close to unity.

The prohibition against copying is problematic. The *no-cloning theorem* [133] explains why it exists. That is, suppose there exists a unitary operator U such that $U(|a\rangle|0\rangle) = |a\rangle|a\rangle$ for all $|a\rangle$. Then, if we have orthogonal states $|a\rangle|0\rangle, |b\rangle|0\rangle$, it follows that (ignoring normalizations) $U(|a\rangle|0\rangle) + U(|b\rangle|0\rangle) = |a\rangle|a\rangle + |b\rangle|b\rangle$. But $U(|a\rangle|0\rangle) + U(|b\rangle|0\rangle) = U(|a\rangle|0\rangle + |b\rangle|0\rangle) = U((|a\rangle + |b\rangle)|0\rangle) = (|a\rangle + |b\rangle)(|a\rangle + |b\rangle) \neq |a\rangle|a\rangle + |b\rangle|b\rangle$.

In some other respects, however, quantum information offers advantages over classical information: by sharing an EPR pair for each transmitted qubit the classical Shannon entropy bound (though not the von Neumann entropy bound) can be violated by up to a factor of 2 (less for submaximal entanglement); this phenomenon is called *superdense coding*. The basic idea [17] is that Alice encodes a two-bit number by applying one of the (four) Pauli matrices to her half of an EPR pair and sending the resulting state to Bob. By performing an XOR to the entangled pair Bob disentangles it; Bob can then perform a measurement on the second qubit (which is then either $|0\rangle$ or $|1\rangle$); Bob then applies the Walsh-Hadamard transform to the first qubit (which is then also either $|0\rangle$ or $|1\rangle$). As we shall see below, superdense coding is related to a dual protocol called *quantum teleportation*.

B. QUANTUM ERROR-CORRECTING CODES

Given that codes exist that allow the faithful maintenance or transmission of quantum information, it is natural to ask what form such a code might take. Indeed, the situation is much the same in the quantum as in the classical regime in that the noiseless coding theorems do not specify a coding scheme. However, the partial correspondence between classical and quantum information provides a partial answer.

For example, although the triple repetition code cannot be carried over into the quantum regime wholesale [because of the no-cloning theorem, e.g., although a (3, 1) quantum quasi-code is allowed which can correct a restricted class of (Boolean) errors], a (9, 1) quantum code based on it exists. Ignoring normalizations,

$$\begin{aligned} |0\rangle \rightarrow |0\rangle_9 &= (|000\rangle + |111\rangle)(|000\rangle + |111\rangle)(|000\rangle + |111\rangle) \\ |1\rangle \rightarrow |1\rangle_9 &= (|000\rangle - |111\rangle)(|000\rangle - |111\rangle)(|000\rangle - |111\rangle) , \end{aligned}$$

and if we consider the (standard) error basis of Pauli matrices,

$$Id, X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} = iXZ ,$$

then the nine-qubit code can correct any linear combination of the errors in the error basis applied to a single qubit [121]. That is, using the parity or majority rule, the first and third terms in the product can correct a Z error; the second term can correct an X error, and hence any one-qubit or Pauli error in the group generated by the Pauli matrices can be corrected. The correction technique uses an ancilla qubit that will, under measurement, collapse its product with a nine-qubit codeword into either an error state with an error recorded in the ancilla (i.e., a syndrome) or into the original state without a recorded error. In either case, measuring the ancilla allows the determination of a single error which can then be explicitly reversed [66].

Calderbank and Shor [34] and Steane [127] first developed the notion of a general *quantum* (n, k) *error-correcting code*, which can be defined as a linear subspace $C = C^k \equiv \mathbb{H}^k$ of \mathbb{H}^n . Their initial work introduced the *CSS codes* constructed from two classical error-correcting codes with $0 \subset C_2 \subset C_1 \subset \mathbb{Z}_2^n$ which respectively correct phase and bit-flip errors. The basic idea is as follows: a CSS code is formed from the derived states (ignoring normalizations)

$$|s_v\rangle \equiv \sum_{w \in C_2} |v + w\rangle ,$$

where $v \in C_1$. It happens that a CSS code is a t -error correcting $(n, \dim C_1 - \dim C_2)$ quantum code, where t is the smaller of the weights of C_1, C_2 . Under the change of basis induced by a Walsh-Hadamard transform, the code maps to the dual code with $0 \subset C_1^\perp \subset C_2^\perp \subset \mathbb{Z}_2^n$, and the codewords themselves map as

$$|s_v\rangle \equiv \sum_{w \in C_2} |v + w\rangle \xrightarrow{W_n} |c_v\rangle \equiv \sum_{u \in C_1} (-1)^{uv} |u\rangle .$$

Phase errors in the original basis map to bit-flip errors in the Walsh-Hadamard basis and vice versa. It therefore suffices to correct bit-flip errors (which are the analogues of classical errors) and perform Walsh-Hadamard transformations, then finally correct any remaining bit-flip errors.

Along with several quantum codes, bounds on quantum code parameters have also been obtained (which have in turn led to the discovery of quantum codes). In [34] a lower bound acting as counterpart to the Levitin-Holevo bound on the asymptotic rates of certain perfect (i.e., having fidelity 1) quantum (n, k) t -error correcting codes was derived: $k/n = 1 - 2H_2(2t/n)$, where $-H_2(x) = x \log_2 x + (1-x) \log_2 (1-x)$. Moreover, a perfect quantum (n, k) t -error correcting code must satisfy $n \geq 4t + k$ [83]. In particular a $(3, 1)$ perfect quantum code that corrects one error violates this bound and hence cannot exist. However, a perfect $(5, 1)$ one-error correcting quantum code exists [89]:

$$\begin{aligned} |0\rangle \rightarrow |0\rangle_5 &= -|00000\rangle + |01111\rangle - |10011\rangle + |11100\rangle + |00110\rangle + |01001\rangle + |10101\rangle + |11010\rangle \\ |1\rangle \rightarrow |1\rangle_5 &= -|11111\rangle + |10000\rangle + |01100\rangle - |00011\rangle + |11001\rangle + |10110\rangle - |01010\rangle - |00101\rangle. \end{aligned}$$

The decoding quantum circuit for the five-qubit code is simply the reversed-order encoding circuit. In this respect the five-qubit code is significant in that it saturates an algebraic bound and can be realized with a minimum of overhead.

If now we set

$$E(J) \equiv \frac{2^k}{2^n} \sum_{j=0}^J 3^j \binom{n}{j},$$

then the *quantum Gilbert-Varshamov lower bound* is $E(d-1) = E(2t) \leq 1$, which reduces in the limit of large n with k/n , t/n fixed to $k/n \leq 1 - \log_2 3 \cdot 2t/n - H_2(2t/n)$. It can be shown that if the Gilbert-Varshamov bound holds, then there exists a quantum (n, k) t -error correcting code [35].

The *quantum Hamming bound* (the analog of the classical Hamming or sphere-packing bound [112] obtained by considering rather than one classical one-bit error three distinct possible single-qubit or Pauli errors) is $E(t) \leq 1$. Although the quantum Hamming bound holds for nondegenerate codes (those codes for which all errors are distinguishable from one another), it is not known whether it applies more generally. In [64], Gottesman showed that a class of $(2^j, 2^j - j - 2)$ 1-error correcting quantum codes saturating the quantum Hamming bound exists.

Quantum error correction can be placed under the broader umbrella of *fault-tolerant quantum computing*; this also encompasses roles such as robust gate application

in the presence of errors (the requirement for which was hinted at in our previous discussion of quantum channels) and entanglement purification or distillation. The interested reader is referred to [20], [65], [113]. Other quantum codes (e.g., quantum stabilizer codes) have been developed; some are touched on in [35], [65].

C. QUANTUM KEY DISTRIBUTION

Before serious consideration was ever given to quantum computation, quantum cryptography was being explored as an instance of the power of quantum information: asymmetric/public-key schemes such as RSA suffer from an inherent susceptibility to computational attack (Appendix A), whereas symmetric/private-key schemes such as DES or IDEA suffer from this *and* the key distribution problem: if the key to a symmetric cipher could already be securely transmitted, then there would be no point in actually transmitting it [99]. Of course, in practice what is done is generally either to distribute keys locally and transport them securely or to encipher keys using an asymmetric scheme (as is the case with PGP). But neither one of these procedures is invulnerable.

Furthermore, there is only one totally secure classical cryptologic protocol: the *one-time pad* or *Vernam cipher* [102]. The scheme is trivial to describe. Let Alice and Bob alone share a perfectly random bitstring. To encrypt a message, Alice simply XORs it with the random bitstring; to decrypt the message Bob does exactly the same. The one-time pad is in fact used where absolute security is paramount; despite its simplicity, however, it is extremely difficult to implement in practice. For example, the key distribution problem is critical, and security demands that only physically secure and authenticated key distribution is acceptable. Moreover, a one-time pad has its name for a reason: using the same pad to encrypt two messages utterly compromises its security. This, coupled with any but the smallest traffic volumes, immediately renders one-time pads infeasible for most practical cryptologic applications [102].

Quantum channels provide a way to make an end run around these problems. Basically, if Eve were to attempt to intercept a quantum key transmission between Alice and Bob, she would inevitably alter the key—and Alice and Bob can use a protocol which exploits this property and subjects the transmitted and received keys to joint statistical tests which establish the security of the transmission. Hence, *quantum key distribution* (QKD) provides security based on the laws of physics rather than the supposed computational infeasibility of inverting one-way functions or of exhaustively searching the keyspace of a cryptosystem (which would be an ideal problem for a quantum computer).

We sketch the basics of the BB84 QKD protocol [16] here. Let $|-\rangle, |+\rangle$ denote the images of $|0\rangle, |1\rangle$ (with $|-\rangle, |0\rangle$ both corresponding to a classical 0 and $|+\rangle, |1\rangle$ both corresponding to a classical 1) under the Walsh-Hadamard transform: both the signed and numbered key pairs form bases (denoted by R and S, respectively) for a one-qubit state space.

To send a key, Alice and Bob perform the following sequence of operations for each bit to be transmitted: Alice first chooses either the R or S basis at random and transmits the state corresponding to a random classical bit in her chosen basis. Bob also chooses one of the bases R, S at random—*independently of Alice*—and performs a measurement. If Alice and Bob used the same basis (and the bit was not intercepted by Eve) then the state encoding the classical random bit will have been perfectly transmitted. With this in mind, Alice and Bob publicly announce their basis selections after all the transmissions are complete: statistically, half of these will agree, and the corresponding classical bits form their provisional shared secret key. To establish its security, Alice and Bob now publicly announce some of the bits of their shared key (which reduces the key length): if these pass certain public statistical procedures (such as testing for a sufficiently low error rate and subsequent classical error correction) then they conclude that the key is secure, since if Eve intercepted the transmission and resent identical states after measuring them in one of the bases, she would send, on average, half her qubits in the wrong basis—and the resultant statistical anomaly would be detected by Alice and Bob. Finally a *privacy amplification* protocol is performed whereby m bitstrings of length n equal to the key length (with $m < n$) are published and the m parities of the XORS are retained as the final key [23].

There are attacks on QKD protocols other than the tapping attack described above. In particular, the *entanglement attack*, in which Eve entangles her interception apparatus with Alice and Bob's quantum channel, is problematic. Similarly, the *swap attack*—whereby Eve stores Alice's transmitted quantum states and sends her own random states to Bob—will succeed occasionally, albeit with exponentially small probability. Of course, any cryptosystem would “suffer” from the exponentially small probability of an adversary correctly guessing a key, and so the security of a QKD protocol is really contingent only on ensuring that the joint probability of the security test passing and Eve gaining more than an exponentially small amount of information about the key is itself exponentially small.

It has recently been shown that several QKD protocols are unconditionally secure in this sense through their relationship to a derived protocol, based on CSS codes, which

is provably secure [123]. The proof furthers work in [98] and [23], wherein the unconditional security of QKD protocols was demonstrated along more complicated lines. Furthermore, it has been shown that error rates of up to 7.56 percent can give asymptotic security (with higher error rates possible under protocol modifications or practical security parameters) [23].

QKD and quantum error correction protocols can also depend on *entanglement distillation or purification* schemes (see below). It is noteworthy that QKD has recently been experimentally realized over 48 km fibers [77] and in daylight over 1.6 km [33].

D. EXPLOITING ENTANGLEMENT: QUANTUM TELEPORTATION AND COMMUNICATION COMPLEXITY

It has been known for some time that transmitting quantum information need not take place over quantum channels; indeed, the protocol of *quantum teleportation* provides an avenue for constructing quantum states from classical information [18]. We omit normalizations throughout the following simplified outline, which otherwise follows [114]. If Alice has a qubit $|\phi\rangle = a|0\rangle + b|1\rangle$ which has not been measured (so that she has no knowledge of a, b) and she and Bob share an EPR pair $|00\rangle + |11\rangle$ (of which, say, the first qubit is Alice's), then Alice considers the state $|\phi\rangle(|00\rangle + |11\rangle) = a|000\rangle + a|011\rangle + b|100\rangle + b|111\rangle$. If now Alice applies a XOR to the first two qubits and then a Walsh-Hadamard transform on the first qubit, as in superdense (de)coding, the three-qubit state becomes

$$|00\rangle(a|0\rangle + b|1\rangle) + |11\rangle(a|1\rangle + b|0\rangle) + |10\rangle(a|0\rangle - b|1\rangle) + |11\rangle(a|1\rangle - b|0\rangle) .$$

A measurement on the first pair collapses the pair, the outcome of which Alice sends to Bob classically. Bob then applies the Pauli matrix corresponding to the result (Id for $|00\rangle$, X for $|01\rangle$, Y for $|10\rangle$ or Z for $|11\rangle$). In the end, Bob is left with the state $|\phi\rangle$, and Alice, in accordance with the no-cloning theorem, is left with a known state which does not depend on $|\phi\rangle$.

As the case is put in [18],

This would appear to offer a more elegant means of private communications than previous quantum cryptographic schemes [BB84, etc.] which require the users to publicly test some of the data exchanged through the quantum channel, in order to certify the privacy of the rest. However, the appearance of intrinsic security is illusory, since an active adversary could effectively tap into the channel by intercepting all the particles on their way to and from Bob, substituting others in such a way as to impersonate

Alice to Bob and Bob to Alice. To defend against this attack Alice and Bob would also need to publicly test some of their data, rendering the present scheme cryptographically equivalent to previous schemes, while retaining its distinctive quantum information-theoretic feature of packing two bits into a single transmitted two-state particle.

Nevertheless, teleportation has a distinct advantage over QKD protocols, for which Alice has to transmit quantum states at the time she wishes to send Bob any information. For teleportation, on the other hand, the requirement is the advance distribution and storage of entangled pairs. Though this is a technical obstacle, overcoming it (a necessary step anyway for a realistic quantum computer) would tilt the balance of utility decisively towards teleportation. More generally, it could be said that teleportation has this intrinsic advantage over any other quantum communication scheme, not least because it avoids the problems of time-of-transmission errors and therefore allows high-fidelity quantum communication.

Indeed, it has been shown [19] that a collection of shared impure entangled pairs can be *distilled* or *purified* into a smaller collection of asymptotically pure entangled pairs, which can in turn be used for faithful teleportation. In this setting the fidelity is replaced with the yield of pure pairs. The purification protocol requires only simple quantum gates (notably, a *bilateral XOR* performed by Alice and Bob on two entangled pairs).

So entanglement can serve as a surrogate for communicating *unknown* information; therefore, it is natural to expect that it can be likewise be exploited in the realm of distributed quantum computation (where after all we are dealing with states which we cannot access without a measurement). Indeed this is the case: Grover [68] showed that a distributed set of coupled EPR pairs could act in parallel to compute the mean of a function (see also V.A) with minimal (classical) communication complexity.

Finally, any “black-box” quantum algorithm (in particular, any amplitude amplification) can be efficiently realized as a related communication protocol. In particular, search algorithms can be realized as quantum communication protocols which offer the same quadratic speedup. This is the well-known “appointment scheduling” result [32].

VII. CONCLUSION

Even considering that this paper is a survey of quantum algorithms and protocols, we have not touched on several elements of the general theory. As far as specifics are concerned, we have deliberately omitted discussion of recent protocols for clock synchronization [38], [81] because these merit in-depth study in their own right. More generally, we have avoided the topics of physical systems for quantum information devices (see, e.g., [79], [128]) and of the impact of decoherence and errors on implementations of protocols and algorithms on quantum information devices (see, e.g., [37], [129]). While strictly speaking neither is premature to address, both of these areas require extraordinary intellectual overhead, and indeed these are precursors to the most important area of concern: applications. Similarly, addressing the problem of spelling out quantum algorithms in circuit models as precursors to their actual implementation is a large undertaking. Some basic scenarios have been examined, however. For example, a proof-of-principle factorization of 15 could be performed with an ion trap quantum computer using as few as 6 qubits and 38 laser pulses [10].

The questions of exploiting Aharonov-Bohm effects for error correction and of how to quantify multipartite entanglement are also of considerable theoretical interest and are as of yet unanswered; these and like issues merit further analysis.

Certainly, if scalable quantum computers are built, then cryptography as we know it is dead. More generally, it is reasonable to assume that a quantum computer would lead to revolutions in physical simulation with the potential to transfigure nanotechnology. Other possible benefits merit consideration also. Problems in combinatorial analysis and statistical decision theory [75], [78] are natural candidates for solution on a quantum computer. Further and as-yet undiscovered applications surely exist.

At this point it is appropriate to say some words about whether a scalable quantum computer will in fact ever be built. The author believes that such devices could well be built within 20 years. Regardless of whether this turns out to be the case, it is certain that current experimental research on manipulating quantum systems will yield dividends. Quantum information has the potential to reshape the world—we have seen why—and at this point it is important to begin considering how.

APPENDIX A

HARD NUMBER-THEORETIC PROBLEMS

APPENDIX A

HARD NUMBER-THEORETIC PROBLEMS

A. THE TROUBLE WITH FACTORING

The difficulty of the factoring problem has long been known. Eratosthenes of Kyrene (ca. 250 BC) provided the first factoring algorithm: given a composite number N , proceed with trial division by all prime numbers less than or equal to its square root.

The *prime number theorem* states that the number of primes less than or equal to N is asymptotically $\pi(N) \approx N/\log N$ [100], and so the sieve of Eratosthenes takes asymptotically as many as $\sqrt{N} \log 2 / \log N$ trial divisions if we have a precomputed list of prime numbers (the construction of which would presumably also require sieving). Thus, factoring a 1,024-bit integer via this method requires on the order of 2^{502} trial divisions: rather a lot. Moreover, it should be remembered that division is computationally expensive.

Over the millennia various advances in factoring have taken place. The current champion of factoring algorithms is the *general number field sieve* (GNFS) [92], of which we provide a technical sketch based on the discussion in [31]. The overall aim (and computationally intensive part, which we will *not* sketch) of the GNFS is to efficiently construct (using a root θ of a monic polynomial f with integer coefficients) a *factor base* U in the number field $\mathbf{Q}(\theta)$ (obtained by adjoining θ to the rational numbers and considering the field that is generated as a result) consisting of algebraic integers (i.e., roots of monic polynomials with integer coefficients also lying in $\mathbf{Q}(\theta)$) [54].

Given such a factor base U such that the product of its elements,

$$\prod_{r=a+b\theta \in U} (a+b\theta) = \alpha^2,$$

is a square of an element of the ring $\mathbf{Z}(\theta)$ generated by θ and

$$\prod (a+bm) = c^2$$

for integers a, b, c , and m with $f(m) \equiv 0 \pmod{N}$, it follows that if we define a subjective ring homomorphism $\phi_m : \mathbf{Z}(\theta) \rightarrow \mathbf{Z}_N$ satisfying $\phi_m(1) = 1$, $\phi_m(\theta) = m$, then

$$\begin{aligned}
x^2 &= \phi_m(f'(\theta)\alpha)^2 = \phi_m\left((f'(\theta)\alpha)^2\right) = \phi_m(f'(\theta))^2 \phi_m(\Pi(a+b\theta)) \\
&\equiv (f'(m))^2 \Pi(a+bm) \equiv y^2 \pmod{N},
\end{aligned}$$

and combining this with the general “difference of squares” result that $\gcd(x \pm y, N)$ divides N for $x^2 \equiv y^2 \pmod{N}$, $x \neq y$, we arrive at a factorization.

The GNFS has asymptotic running time $L_N[1/3, (8/3)^{2/3}]$, where

$$L_N[\alpha, c] = O\left(\exp\left((c + o(1))(\log N)^\alpha (\log \log N)^{1-\alpha}\right)\right)$$

is the *Lucas complexity class* [102]. Using this, we arrive at estimates of operation counts required for the GNFS. If we use as a benchmark 1,000 Macintosh G4 computers running the GNFS totally in parallel (the GNFS is highly parallelizable) at 1 Gflops each and identify this with 1 G(GNFS)ops, then the corresponding operation counts and runtime order estimates are as follows:

Bit length	$O(\text{Ops})$	$O(\text{time})$ (s)	$O(\text{time})$ (yr)
512	2^{64}	$1.8 \cdot 10^7$	5.6
1,024	2^{87}	$1.3 \cdot 10^{14}$	$4.2 \cdot 10^7$
2,048	2^{117}	$1.5 \cdot 10^{23}$	$4.2 \cdot 10^{16}$

Hence, factoring, for example, a 2,048-bit integer is computationally infeasible using the GNFS.

A hard 512-bit number (RSA-155) was recently (August 1999) factored via a massively distributed sieving effort using roughly 300 computers—over half of them high-end workstations—over the course of 7.4 months (5.2 months for sieving and 2.2 months to select an appropriate monic polynomial for the GNFS), plus the final solution time of the resulting massive sparse linear system via a specialized iterative technique [31], which required roughly 10 days on a Cray C916 [130].

It is therefore reasonable to assume that factoring 512-bit numbers is well within a temporal-computational scope corresponding to a capital outlay on the order of \$10 million and a year of execution time for GNFS or a slightly better algorithm. Under these assumptions, however, 1,024-bit numbers are still inaccessible—at any price. Silverman notes in an RSA technical report that the bit length of the largest number openly factored as a function of the year has a linear fit ($b = 14.05[y - 1970] + 23$) with correlation coefficient .955, where b is the bit length and y is the year [124]. Another estimate based on extrapolating Moore’s law arrives at a cube root fit which is cited in

[124]. Assuming these relationships hold indefinitely, we get the following estimates of the year of factorization capability for a given bit length:

Bit length	Linear fit	Moore's law
512	2,005	1,999
1,024	2,041	2,018
2,048	2,115	2,041

This suggests that to factor numbers much larger than 512 bits it is better to wait for developments in algorithms and computers than to bother with the GNFS.

B. RSA

Factoring is not just an academic exercise; indeed, the security of the nearly universal standard RSA public-key cryptosystem [115] hinges on the computational infeasibility of factoring large numbers. We present a sketch of the number-theoretic problem upon which the RSA protocols are based.

Alice puts an *RSA modulus* $N = pq$ for two large (and otherwise suitable) prime numbers p, q of equal or nearly equal length and picks an *encryption key* e such that

$$\gcd(e, (p-1)(q-1)) = 1.$$

She can efficiently compute the *decryption key*

$$d = e^{-1} \bmod (p-1)(q-1).$$

Alice publishes N and e , and keeps d, p , and q secret. If Bob wishes to send Alice a secret message M (here, just a number less than N), he encrypts it as $C = M^e \bmod N$. Alice then computes

$$C^d \bmod N \equiv M^{de} \bmod N \equiv M^{k(p-1)(q-1)+1} \bmod N \equiv M^{k\phi(N)+1} \bmod N \equiv M \bmod N,$$

where the *Euler phi function* $\phi(N)$ is defined as the number of positive integers less than and relatively prime to N (in our case equal to $(p-1)(q-1)$), and we have invoked Euler's theorem [88]:

$$\gcd(a, N) = 1 \Rightarrow a^{\phi(N)} \equiv 1 \bmod N.$$

Since (by assumption and design) M is less than N , we recover the message uniquely. (N.B. By decomposing an arbitrary message into packets we can always do this.)

The *RSA problem* is to derive M given N, e , and C ; it is generally suspected [102] (though not known) that this is polytime equivalent to factoring (certainly factoring

moduli gives efficient solutions of RSA). Therefore, we may reinterpret the tables from the previous section as security parameters for the RSA cryptosystem, and provide a new context in which the factoring problem may be said to be important.

C. THE DISCRETE LOGARITHM PROBLEM

If p is a prime number, the multiplicative group \mathbf{Z}_p^* is cyclic, and the *discrete logarithm problem* (DLP) for a generator α and arbitrary unit β is to determine (the unique) x such that $\alpha^x \equiv \beta \pmod{p}$ [102]. The DLP generalizes to algebraic curves with group structures [87], but we shall not consider these here.

The ElGamal [55] and Digital Signature Algorithm [59] schemes (among others) rely on the DLP. Though ElGamal can be used for encryption, we sketch here only the basis for the authentication protocol. In this setting, Alice randomly picks two elements α and x in the cyclic group \mathbf{Z}_p^* and computes

$$\beta = \alpha^x \pmod{p}.$$

Alice publishes p , α and β and keeps x secret. To authenticate a message M , Alice chooses a secret *signature exponent* k and computes

$$a = \alpha^k \pmod{p}, \text{ } b \text{ such that } M = (ax + bk) \pmod{p-1}.$$

The public pair a, b is the *signature*. Authentication proceeds along the following lines:

$$\beta^a a^b \pmod{p} = \alpha^{ax} \alpha^{bk} \pmod{p} = \alpha^{ax+bk} \pmod{p} = \alpha^M \pmod{p}.$$

DSA is similar in its operation, and in fact it can be shown [118] that they are both cases of a general DLP signature scheme for cyclic groups.

Shor also provided in [122] an quantum-algorithmic solution to the DLP along much the same lines (and with a basically equal increase in efficiency) as for the factoring problem; Boneh and Lipton [27] obtained an analogous for algebraic curves. Kitaev's solution of the ASP encompasses these [83].

APPENDIX B

CLASSICAL INFORMATION THEORY

APPENDIX B

CLASSICAL INFORMATION THEORY

A. CLASSICAL ENTROPY, CHANNEL CAPACITY, AND ERROR CORRECTION

Shannon [120] initiated the study of information theory; its basic building blocks are the notions of *entropy* and *channel capacity*. (The interested reader may also refer to [5] or [99] for brief or detailed discussions, respectively.) Given a statistical characterization of a discrete channel—that is, given a random variable X which takes as its values the possible transmissions or events E_1, \dots, E_n and their (presumably nonzero) associated probabilities p_1, \dots, p_n , a reasonable measure I of information transmitted should satisfy the following criteria:

- I. $I\left(\bigcap_j E_j\right) \geq \max_j I(E_j) \geq 0$
- II. $I\left(\bigcap_j E_j\right) = \sum_j I(E_j)$ for independent events.

It can then be shown that I must be of the form $I(E_j) = -\log_2 p_j$ (up to a multiplicative constant) and so its expected value—the *entropy*—is

$$H(X) = \langle I(X) \rangle = -\sum_{j=1}^n p_j \log_2 p_j.$$

In this context, the entropy can be said to be the appropriate measure of information (properly, of uncertainty) which is transmitted through a communication channel. We may also define the respective *joint and conditional entropies* for X, Y by

$$H(X, Y) = -\sum_{j=1}^n \sum_{k=1}^m p_{jk} \log_2 p_{jk}$$

$$H_X(Y) = -\sum_{j=1}^n p_j \left(\sum_{k=1}^m p_j(k) \log_2 p_j(k) \right) = \sum_{j=1}^n p_j H_j(Y),$$

where $p_j(k) = p(Y=k|X=j)$ is a conditional probability and the sum in parentheses is called the *equivocation*. It turns out that $H(X, Y) = H(X) + H_X(Y) = H(Y) + H_Y(X)$. Finally, we define the *mutual information* (Shannon's *transmission rate*) $M(X, Y) = H(X) - H_Y(X) = H(Y) - H_X(Y)$: since the conditional entropy is a measure of residual information, the mutual information is what it should be.

The *channel capacity* C is then the maximum possible value of the mutual information. Shannon proved the following theorem:

Let a discrete channel have the capacity C and a discrete source the entropy per second H . If $H \leq C$ there exists a coding system such that the output of the source can be transmitted over the channel with an arbitrarily small frequency of errors (or an arbitrarily small equivocation). If $H > C$ it is possible to encode the source such that the equivocation is less than $H - C + \epsilon$ where ϵ is arbitrarily small. There is no method of encoding which gives an equivocation less than $H - C$.

Establishing the existence of good error-correcting codes is, however, a far cry from having (or being able to implement) good error-correcting codes.

The simplest example of an error-correcting code is the *triplet parity code*: 0 is encoded as the codeword 000 and 1 as 111. A received triplet other than these is weighted: either it has two zeroes or two ones, according to which it is changed to 000 or 111 accordingly. This is a specific instance $(3, 1)$ of the more general notion of a *linear binary (n, k) or (n, k, d) code*. (Here, d refers to the minimum weight, or number of ones, in a codeword, and it can be shown that an (n, k, d) code can correct $(d-1)/2$ or fewer errors; the [integral] number t of errors a code can correct is referred to as its *weight*.) Such a code C is specified by, for example, a *generator matrix* G which can be assumed to be in the form $(Id|A)$, where Id is the k -by- k identity matrix and A is a k -by- $(n-k)$ matrix (equivalently, the *dual code* C^\perp may be characterized by the *parity check matrix* $(-A^T|Id)$). The rows of the matrix G are then the basis codewords, and a generic bit string x of length k is encoded by producing the linear combination of basis codewords whose first k bits equal x . Hence, a linear code can also be described by the span of its basis codewords; this turns out to be the view most naturally suited to negotiating the correspondence between classical and quantum codes.

The decoding process is generally difficult: each codeword has a large co-set of errorwords which (unless the code were engineered with viable algorithmic decoding schemes) has to be exhaustively searched. However, special decoding techniques exist (e.g., syndrome and Hamming decoding) which can dramatically reduce the computational effort involved. Still, when n is large enough, an (n, k) code is infeasible to implement classically (if for no other reason than that processing with such a code is problematic from the standpoint of buffer size, bus speed, etc.).

It turns out [112] that the $(2^{n-k}-1, k, 3)$ Hamming and $(23, 12, 7)$ Golay codes are the only nontrivial binary perfect (i.e., capable of correcting t errors) error-correcting

codes. This surprising fact serves to illustrate that the theory of classical error-correcting codes is deep and complex. We refer the reader to [111] or [112] for further details.

APPENDIX C

REFERENCES

APPENDIX C

REFERENCES

- [1] Abrams, D.S., and Lloyd, S. "Simulation of Many-Body Fermi Systems on a Universal Quantum Computer." *Phys. Rev. Lett.* **79**, 2586 (1997).
- [2] ———. "A Quantum Algorithm Providing Exponential Speed Increase for Finding Eigenvalues and Eigenvectors." quant-ph/9807070 (1998).
- [3] Abrams, D.S., and Williams, C.P. "Fast Quantum Algorithms for Numerical Integrals and Stochastic Processes." quant-ph/9908083 (1999).
- [4] Aharonov, D. "Quantum Computation." quant-ph/9812037 (1998).
- [5] Applebaum, D. *Probability and Information: an integrated approach*. Cambridge University Press, Cambridge (1996).
- [6] Barenco, A., Bennett, C.H., Cleve, R., DiVincenzo, D.P., Margolus, N., Shor, P., Sleator, T., Smolin, J.A., and Weinfurter, H. "Elementary gates for quantum computation." *Phys. Rev. A* **52**, 3457 (1995).
- [7] Barenco, A., Ekert, A., Suominen, K., and Törmä, P. "Approximate quantum Fourier transform and decoherence." *Phys. Rev. A* **54**, 139 (1996).
- [8] Balko, B. Personal communication (2000).
- [9] Beals, R., Buhrman, H., Cleve, R., Mosca, M., and de Wolf, R. "Quantum Lower Bounds by Polynomials." quant-ph/9802049 (1998).
- [10] Beckman, D., Chari, A. N., Devabhaktuni, S., and Preskill, J. "Efficient networks for quantum factoring." *Phys. Rev. A* **54**, 1034 (1996).
- [11] Beckman, F. S. *Mathematical Foundations of Programming*. Addison-Wesley, Reading, Massachusetts (1980).
- [12] Bell, J. S. *Physics* **1**, 195 (1964).
- [13] Benioff, P. "The computer as a physical system: a microscopic quantum mechanical Hamiltonian model of computers as represented by Turing machines." *J. Stat. Phys.* **22**, 563 (1980).
- [14] Benjamin, S.C., and Johnson, N.F. "Cellular Structures for Computation in the Quantum Regime." cond-mat/9808243 (1998).
- [15] Bennett, C.H. "Logical reversibility of computation." *IBM J. Res. Dev.* **17**, 525 (1973).
- [16] Bennett, C.H., and Brassard, G., in *Proceedings of the IEEE International Conference on Computers, Systems and Signal Processing, Bangalore, India*. IEEE, New York (1984).

- [17] Bennett, C.H., and Wiesner, S.J. "Communication via One- and Two-Particle Operators on Einstein-Podolsky-Rosen States." *Phys. Rev. Lett.* **69**, 2881 (1992).
- [18] Bennett, C.H., Brassard, G., Crépeau, C., Josza, R., Peres, A., and Wootters, W. "Teleporting an Unknown Quantum State via Dual Classical and Einstein-Podolsky-Rosen Channels." *Phys. Rev. Lett.* **70**, 1895 (1993).
- [19] Bennett, C.H., Brassard, G., Popescu, S., Schumacher, B., Smolin, J.A., and Wootters, W. "Purification of Noisy Entanglement and Faithful Teleportation via Noisy Channels." *Phys. Rev. Lett.* **76**, 722 (1996).
- [20] Bennett, C.H., DiVincenzo, D.P., Smolin, J.A., and Wootters, W.K. "Mixed-state entanglement and quantum error correction." *Phys. Rev. A* **54**, 3824 (1996).
- [21] Bennett, C.H., and DiVincenzo, D.P. "Quantum information and computation." *Nature* **404**, 247 (2000).
- [22] Bernstein, E., and Vazirani, U. in *Proceedings of the 25th Annual ACM Symposium on Theory of Computing*, ACM, New York (1993).
- [23] Biham, E., Boyer, M., Boykin, P. O., Mor, T., and Roychowdhury, V. "A Proof of the Security of Quantum Key Distribution." quant-ph/9912053 (1999).
- [24] Biroli, G., Monasson, R., and Weight, M. "A variational description of the ground state structure in random satisfiability problems." cond-mat/9907343 (1997).
- [25] Boghosian, B.M., and Taylor, Washington IV. "A Quantum Lattice-Gas Model for the Many-Particle Schrödinger Equation in d Dimensions." quant-ph/9604035 (1997).
- [26] Bohm, A. *Quantum Mechanics—Foundations and Applications*. 3rd ed. Springer, New York (1993).
- [27] Boneh, D., and Lipton, R.J., in *Advances in Cryptology—CRYPTO '95*. Springer, Berlin (1995).
- [28] Boyer, M., Brassard, G., Høyer, P., and Tapp, A. "Thigh Bounds on Quantum Searching." quant-ph/9605034 (1996).
- [29] Brassard, G., Høyer, P., and Tapp, A. "Quantum Counting." quant-ph/9805082 (1998).
- [30] Braunstein, S.L. "Quantum Computation." Preprint (1998).
- [31] Briggs, M.E. "An Introduction to the General Number Field Sieve." Master's thesis, Virginia Tech (1998).
- [32] Buhrman, H., Cleve, R., and Wigderson, A. "Quantum vs. Classical Communication and Computation." quant-ph/9802040 (1998).
- [33] Buttler, W.T., Hughes, R.J., Lamoreaux, S.K., Morgan, G.L., Nordholt, J.E., and Peterson, C.G. "Daylight Quantum Key Distribution over 1.6 km." *Phys. Rev. Lett.* **84**, 5652 (2000).
- [34] Calderbank, A.R., and Shor, P.W. "Good quantum error-correcting codes exist." *Phys. Rev. A* **54**, 1098 (1996).

- [35] Calderbank, A.R., Rains, E.M., Shor, P.W., and Sloane, N.J.A. "Quantum Error Correction and Orthogonal Geometry." *Phys. Rev. Lett.* **78**, 405 (1997).
- [36] Cerf, N.J., Grover, L.K., and Williams, C.P. "Nested quantum search and structured problems." *Phys. Rev. A* **61**, 032303 (2000).
- [37] Chuang, I.L., Laflamme, R., Shor, P., and Zurek, W.H. "Quantum Computers, Factoring, and Decoherence." quant-ph/9503007 (1995).
- [38] Chuang, I.L. "Quantum Algorithm for Distributed Clock Synchronization." quant-ph/0005092 (2000).
- [39] Chopard, B., and Droz, M. *Cellular Automata Modeling of Physical Systems*. Cambridge University Press, Cambridge (1998).
- [40] Cleve, R., and Buhrman, H. "Substituting Quantum Entanglement for Communication." quant-ph/9704026 (1997).
- [41] Cleve, R., Ekert, A., Macciavello, C., and Mosca, M. "Quantum Algorithms Revisited." quant-ph/9708016 (1998).
- [42] Cleve, R., and Watrous, J. "Fast Parallel Circuits for the Quantum Fourier Transform." quant-ph/0006004 (2000).
- [43] Conway, J.B. *A Course in Functional Analysis*. 2nd ed. Springer, New York (1990).
- [44] Coppersmith, D. "An approximate Fourier transform useful in quantum factoring." IBM research report RC19642 (1994).
- [45] Deutsch, D. "Quantum theory, the Church-Turing principle and the universal quantum computer." *Proc. Roy. Soc. A* **400**, 97 (1985).
- [46] Deutsch, D., and Josza, R. "Rapid solution of problems by quantum computation." *Proc. Roy. Soc. A* **439**, 553 (1992).
- [47] Deutsch, D. unpublished (1994).
- [48] Deutsch, D., Barenco, A., and Ekert, A. "Universality in quantum computation." *Proc. Roy. Soc. A* **449**, 669 (1995).
- [49] DiVincenzo, D.P. "Quantum Computation." *Science* **270**, 255 (1995).
- [50] ———. "Two-bit gates are universal for quantum computation." *Phys. Rev. A* **51**, 1015 (1995).
- [51] ———. "The Physical Implementation of Quantum Computation." quant-ph/0002077 (2000).
- [52] Dowling, J. Personal communication (2000).
- [53] Draper, T.G. "Addition on a Quantum Computer." NSA/R51 technical report (unclassified, 1998).
- [54] Dummitt, D.S., and Foote, R.M. *Abstract Algebra*. Prentice Hall, Englewood Cliffs, New Jersey (1991).

- [55] ElGamal, T. in *Advances in Cryptology—CRYPTO '84*. Springer, Berlin (1985).
- [56] Feynman, R.P. "Simulating Physics with Computers." *Int. J. Theor. Phys.* **21**, 467 (1982).
- [57] ———. "Quantum Mechanical Computers." *Found. Phys.* **16**, 507 (1986).
- [58] Feynman, R.P., and Hibbs, A.R. *Quantum Mechanics and Path Integrals*. McGraw-Hill, New York (1965).
- [59] Federal Information Processing Standard 186, "Digital signature standard." csrc.nist.gov/fips/fips1861.pdf
- [60] Fredkin, E., and Toffoli, T. "Conservative logic." *Int. J. Theor. Phys.* **21**, 219 (1982).
- [61] Gershenfeld, N. *The Nature of Mathematical Modeling*. Cambridge University Press, Cambridge (1999).
- [62] Gingrich, R.M., Williams, C.P., and Cerf, N.J. "Generalized quantum search with parallelism." *Phys. Rev. A* **61**, 052313 (2000).
- [63] Goldenfeld, N. *Lectures on Phase Transitions and the Renormalization Group*. Addison-Wesley, Reading, Massachusetts (1992).
- [64] Gottesman, D. "Class of quantum error-correcting codes saturating the quantum Hamming bound." *Phys. Rev. A* **54**, 1862 (1996).
- [65] ———. "Theory of fault-tolerant quantum computation." *Phys. Rev. A* **57**, 127 (1998).
- [66] ———. "An Introduction to Quantum Error Correction." quant-ph/0004072 (2000).
- [67] Grover, L.K. "Quantum mechanics helps in searching for a needle in a haystack." *Phys. Rev. Lett.* **79**, 325 (1997).
- [68] ———. "Quantum Telecomputation." quant-ph/9704012 (1997).
- [69] ———. "A Framework for Fast Quantum Mechanical Algorithms." quant-ph/9711043 (1997).
- [70] Hardy, G.H., and Wright, E.M. *An Introduction to the Theory of Numbers*. 5th ed. Oxford University Press, New York (1979).
- [71] Hausladen, P., Josza, R., Schumacher, B., Westmoreland, M., and Wootters, W. K. "Classical information capacity of a quantum channel." *Phys. Rev. A* **54**, 1869 (1996).
- [72] Heagy, J. IDA memorandum (unclassified, 1999).
- [73] Hogg, T. "Highly Structured Searches with Quantum Computers." *Phys. Rev. Lett.* **80**, 2473 (1998).
- [74] ———. "Quantum search heuristics." *Phys. Rev. A* **61**, 052311 (2000).
- [75] Hogg, T., and Portnov, D. "Quantum Optimization." quant-ph/0006090 (2000).

- [76] Høyer, P. "On Arbitrary Phases in Quantum Amplitude Amplification." quant-ph/0006031 (2000).
- [77] Hughes, R.J., Morgan, G.L., and Peterson, C.G. "Practical quantum key distribution over a 48-km optical fiber network." Los Alamos technical report LA-UR-99-1593 (1999).
- [78] Huntsman, S. unpublished (2000).
- [79] Jones, J.A. "NMR Quantum Computation: a Critical Evaluation." quant-ph/0002085 (2000).
- [80] Josza, R. "Quantum Algorithms and the Fourier Transform." quant-ph/9707033 (1997).
- [81] Josza, R., Abrams, D.S., Dowling, J.P., and Williams, C.P. "Quantum Clock Synchronization Based on Shared Prior Entanglement." quant-ph/0004105 (2000).
- [82] Kirkpatrick, S., and Selman, B. "Critical Behavior in the Satisfiability of Random Boolean Expressions." *Science* **264**, 1297 (1994).
- [83] Kitaev, A.Y. "Quantum Measurements and the Abelian Stabilizer Problem." quant-ph/9511026 (1995).
- [84] Knill, E., and Laflamme, R. "Theory of quantum error-correcting codes." *Phys. Rev. A* **55**, 900 (1997).
- [85] Knill, E., Laflamme, R., Martinez, R., and Tseng, C.-H. "An algorithmic benchmark for quantum information processing." *Nature* **404**, 368 (2000).
- [86] Koblitz, N. *A Course in Number Theory and Cryptography*. 2nd ed. Springer, New York (1994).
- [87] ———. *Algebraic Aspects of Cryptography*. Springer, Berlin (1998).
- [88] Knuth, D.E. *The Art of Computer Programming*, vols. 1-3, Addison-Wesley, Reading, Massachusetts (1997).
- [89] Laflamme, R., Miquel, C., Paz, J.P., and Zurek, W. "Perfect Quantum Error Correcting Code." *Phys. Rev. Lett.* **77**, 198 (1996).
- [90] Laméhi-Rachti, M., and Mittag, W. *Phys. Rev. D* **14**, 2543 (1976).
- [91] Landauer, R. "Irreversibility and heat generation in the computing process." *IBM J. Res. Dev.* **3**, 183 (1961).
- [92] Lenstra, A.K., and Lenstra, H.W. (eds.) *The development of the number field sieve*, LNM 1554. Springer, Berlin (1993).
- [93] Lidar, D.A., and Biham, O. "Simulating Ising Spin Glasses on a Quantum Computer." quant-ph/9611038 (1997).
- [94] Lloyd, S. "A Potentially Realizable Quantum Computer." *Science* **261**, 1569 (1993).

- [95] ———. “Almost Any Quantum Logic Gate is Universal.” *Phys. Rev. Lett.* **75**, 346 (1995).
- [96] ———. “Universal Quantum Simulators.” *Science* **273**, 1073 (1996).
- [97] ———. Personal communication (2000).
- [98] Mayers, D. in *Advances in Cryptology—CRYPTO '96*. Springer, Berlin (1996).
- [99] McEliece, R.J. *The Theory of Information and Coding*. Addison-Wesley, Reading, Massachusetts (1977).
- [100] McKean, H.P. *Fourier Series and Integrals*. Academic Press, San Diego (1972).
- [101] Mehra, J. *The Beat of a Different Drum*. Oxford University Press, New York (1994).
- [102] Menezes, A.J., van Oorschot, P.C., and Vanstone, S.A. *Handbook of Applied Cryptography*. CRC Press, Boca Raton (1997).
- [103] Meyer, D.A. “From Quantum Cellular Automata to Quantum Lattice Gases.” quant-ph/9604003 (1996).
- [104] Miller, G.L. “Riemann’s hypothesis and tests for primality.” *J. Comp. Sys. Sci.* **13**, 300 (1976).
- [105] Monasson, R., Zecchina, R., Kirkpatrick, S., Selman, B., and Troyansky, L. “Determining computational complexity from characteristic ‘phase transitions.’” *Nature* **400**, 133 (1999).
- [106] Monroe, C., Meekhof, D.M., King, B.E., and Wineland, D.J. “A ‘Schrödinger cat’ superposition state of an atom.” *Science* **272**, 1131 (1996).
- [107] Myatt, C.J., King, B.E., Turchette, Q.A., Sackett, C.A., Kielpinski, D., Itano, W.M., Monroe, C., and Wineland, D.J. “Decoherence of quantum superpositions through coupling to engineered reservoirs.” *Nature* **403**, 269 (2000).
- [108] Nayak, A., and Wu, F. “The Quantum Query Complexity of Approximating the Median and Related Statistics.” quant-ph/9804066 (1998).
- [109] Omnès, R. *The Interpretation of Quantum Mechanics*. Princeton University Press, Princeton (1994).
- [110] Peskin, M.E., and Schroeder, D.V. *An Introduction to Quantum Field Theory*. Addison-Wesley, Reading, Massachusetts (1995).
- [111] Peterson, W.W., and Weldon, E.J. Jr. *Error-Correcting Codes*. 2nd ed. MIT Press, Cambridge, Massachusetts (1972).
- [112] Pless, V. *Introduction to the Theory of Error-Correcting Codes*. John Wiley, New York (1982).
- [113] Preskill, J. “Fault-Tolerant Quantum Computation.” quant-ph/9712048 (1997).
- [114] Rieffel, E., and Polak, W. “An Introduction to Quantum Computing for Non-Physicists.” quant-ph/9809016 (2000).

- [115] Rivest, R.L., Shamir, A., and Adelman, L.M. "A method for obtaining digital signatures and public-key cryptosystems." *Comm. ACM* **21**, 120 (1978).
- [116] Sackett, C.A., Kielpinski, D., King, B.E., Langer, C., Meyer, V., Myatt, C.J., Rowe, M., Turchette, G.A., Itano, W.M., Wineland, D.J., and Monroe, C. "Experimental entanglement of four particles." *Nature* **404**, 256 (2000).
- [117] Sattinger, D.H., and Weaver, O.L. *Lie Groups and Lie Algebras with Applications to Physics, Geometry, and Mechanics*. Springer, New York (1986).
- [118] Schneier, B. *Applied Cryptography*. John Wiley, New York (1996).
- [119] Schumacher, B. "Quantum coding." *Phys. Rev. A* **51**, 2738 (1995).
- [120] Shannon, C.E., and Weaver, W. *The Mathematical Theory of Communication*. University of Illinois Press, Urbana (1998).
- [121] Shor, P.W. "Scheme for reducing decoherence in quantum computer memory." *Phys. Rev. A* **52**, R2493 (1995).
- [122] ———. "Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer." quant-ph/9508027 (1996).
- [123] Shor, P.W., and Preskill, J. "Simple Proof of Security of the BB84 Quantum Key Distribution Protocol." quant-ph/0003004 (2000).
- [124] Silverman, R.D. "A Cost-Based Security Analysis of Symmetric and Asymmetric Key Lengths." www.rsasecurity.com/rsalabs/bulletins/bulletin13.html (2000).
- [125] Simon, D. in *Proceedings of the 35th Annual Symposium on Foundations of Computer Science*, IEEE, Los Alamitos, California (1994).
- [126] Sleator, T., and Weinfurter, H. "Realizable Universal Quantum Logic Gates." *Phys. Rev. Lett.* **74**, 4087 (1995).
- [127] Steane, A.M. "Error Correcting Codes in Quantum Theory." *Phys. Rev. Lett.* **77**, 793 (1996).
- [128] Steane, A.M., and Lucas, D.M. "Quantum Computing with Trapped Ions, Atoms, and Light." quant-ph/0004053 (2000).
- [129] Unruh, W.G. "Maintaining coherence in quantum computers." *Phys. Rev. A* **51**, 992 (1995).
- [130] www.rsasecurity.com/rsalabs/challenges/factoring/rsa155.html (1999).
- [131] Toffoli, T., and Margolus, N. *Cellular Automata Machines*. MIT Press, Cambridge, Massachusetts (1987).
- [132] Vedral, V., Barenco, A., and Ekert, A. "Quantum networks for elementary arithmetic operations." *Phys. Rev. A* **54**, 147 (1996).
- [133] Wootters, W.K., and Zurek, W.H. "A single quantum cannot be cloned." *Nature* **299**, 802 (1982).
- [134] Yezpez, J. Personal communication (2000).

- [135] ———. in *Quantum Computing and Quantum Communication Lecture Notes in Computer Science*, LNCS 1509. Springer, Berlin (1999).
- [136] Zalka, C. “Grover’s Quantum Searching Algorithm is Optimal.” quant-ph/9711070 (1999).
- [137] Zurek, W.H. “Decoherence and the Transition from Quantum to Classical.” *Physics Today* **44**, 36 (1991).
- [A1] Friedman, J.H., Patel, V., Chen, W., Tolpygo, S.K., and Lukens, J.E., “Quantum superposition of distinct macroscopic states.” *Nature* **406**, 43 (2000).

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public Reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE February 2001	3. REPORT TYPE AND DATES COVERED Final — May — July 2000
4. TITLE AND SUBTITLE Quantum Algorithms and Protocols			5. FUNDING NUMBERS DASW01 98 C 0067 CRP-2048
6. AUTHOR(S) Steve Huntsman			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Institute for Defense Analyses 1801 N. Beauregard St. Alexandria, VA 22311-1772			8. PERFORMING ORGANIZATION REPORT NUMBER IDA Paper P-3540
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSORING/MONITORING AGENCY REPORT NUMBER
11. SUPPLEMENTARY NOTES			
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited.			12b. DISTRIBUTION CODE
13. ABSTRACT (Maximum 180 words) New results from (and attempts to recast physics itself as) information-theoretic interpretations of nature have led to significant progress in physics and in computer science. Foremost among the efforts in this vein is quantum information, which, largely on the basis of startling results on quantum teleportation and polynomial-time factoring, has evolved into a major scientific initiative. Quantum information traces its roots back to the famous Schrödinger's cat <i>gedanken</i> experiment, which highlights the dual natures of quantum information: namely, entanglement and decoherence. The resolution of the cat paradox has significant implications for national security and the physical and informational sciences. Its potentially profound effect on the evolution of these areas is the motivation for our discussion. We present a summary overview of theoretical results in quantum information which quickly covers most of the terrain (at the cost of difficulty). Although fairly comprehensive, gaps in coverage include entanglement measures, decoherence-free spaces/subsystems, and holonomic quantum protocols for computation and error correction. Appendices on classical factoring and information theory are included.			
14. SUBJECT TERMS quantum information, cryptology			15. NUMBER OF PAGES 61
			16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT SAR